

Integer Set Library: Manual

Version: isl-0.03-16-ga047c60

Sven Verdoolaege

July 24, 2010

Contents

| | | |
|----------|---|----------|
| 1 | User Manual | 3 |
| 1.1 | Introduction | 3 |
| 1.1.1 | Backward Incompatible Changes | 3 |
| 1.2 | Installation | 3 |
| 1.2.1 | Installation from the git repository | 4 |
| 1.2.2 | Common installation instructions | 4 |
| 1.3 | Library | 5 |
| 1.3.1 | Initialization | 5 |
| 1.3.2 | Integers | 6 |
| 1.3.3 | Sets and Relations | 8 |
| 1.3.4 | Memory Management | 8 |
| 1.3.5 | Dimension Specifications | 9 |
| 1.3.6 | Input and Output | 10 |
| 1.3.7 | Creating New Sets and Relations | 12 |
| 1.3.8 | Inspecting Sets and Relations | 15 |
| 1.3.9 | Properties | 17 |
| 1.3.10 | Unary Operations | 19 |
| 1.3.11 | Binary Operations | 21 |
| 1.3.12 | Points | 25 |
| 1.3.13 | Piecewise Quasipolynomials | 26 |
| 1.3.14 | Bounds on Piecewise Quasipolynomials and Piecewise Quasipolynomial Reductions | 30 |
| 1.3.15 | Dependence Analysis | 32 |
| 1.3.16 | Parametric Vertex Enumeration | 33 |
| 1.4 | Applications | 35 |
| 1.4.1 | <code>isl_polyhedron_sample</code> | 35 |
| 1.4.2 | <code>isl_pip</code> | 35 |
| 1.4.3 | <code>isl_polyhedron_minimize</code> | 35 |
| 1.4.4 | <code>isl_polytope_scan</code> | 35 |
| 1.5 | <code>isl-polylib</code> | 35 |

| | | |
|----------|--|-----------|
| 2 | Implementation Details | 37 |
| 2.1 | Sets and Relations | 37 |
| 2.2 | Simple Hull | 38 |
| 2.3 | Coalescing | 39 |
| 2.4 | Transitive Closure | 39 |
| 2.4.1 | Introduction | 39 |
| 2.4.2 | Computing an Approximation of R^k | 40 |
| 2.4.3 | Checking Exactness | 42 |
| 2.4.4 | Decomposing R into strongly connected components | 43 |
| 2.4.5 | Partitioning the domains and ranges of R | 46 |
| 2.4.6 | Incremental Computation | 48 |
| 2.4.7 | An Omega-like implementation | 49 |
| | References | 51 |

Chapter 1

User Manual

1.1 Introduction

`isl` is a thread-safe C library for manipulating sets and relations of integer points bounded by affine constraints. The descriptions of the sets and relations may involve both parameters and existentially quantified variables. All computations are performed in exact integer arithmetic using GMP. The `isl` library offers functionality that is similar to that offered by the `Omega` and `Omega+` libraries, but the underlying algorithms are in most cases completely different.

The library is by no means complete and some fairly basic functionality is still missing. Still, even in its current form, the library has been successfully used as a backend polyhedral library for the polyhedral scanner `CLOoG` and as part of an equivalence checker of static affine programs. For bug reports, feature requests and questions, visit the the discussion group at <http://groups.google.com/group/isl-development>.

1.1.1 Backward Incompatible Changes

Changes since `isl-0.02`

- The old printing functions have been deprecated and replaced by `isl_printer` functions, see [Input and Output](#).
- Most functions related to dependence analysis have acquired an extra `must` argument. To obtain the old behavior, this argument should be given the value 1. See [Dependence Analysis](#).

1.2 Installation

The source of `isl` can be obtained either as a tarball or from the git repository. Both are available from <http://freshmeat.net/projects/isl/>. The installation process depends on how you obtained the source.

1.2.1 Installation from the git repository

1. Clone or update the repository

The first time the source is obtained, you need to clone the repository.

```
git clone git://repo.or.cz/isl.git
```

To obtain updates, you need to pull in the latest changes

```
git pull
```

2. Get submodule (optional)

isl can optionally use the `piplib` library and provides this library as a submodule. If you want to use it, then after you have cloned `isl`, you need to grab the submodules

```
git submodule init
git submodule update
```

To obtain updates, you only need

```
git submodule update
```

Note that `isl` currently does not use any `piplib` functionality by default.

3. Generate configure

```
./autogen.sh
```

After performing the above steps, continue with the Common installation instructions.

1.2.2 Common installation instructions

1. Obtain GMP

Building `isl` requires `GMP`, including its headers files. Your distribution may not provide these header files by default and you may need to install a package called `gmp-devel` or something similar. Alternatively, `GMP` can be built from source, available from <http://gmplib.org/>.

2. Configure

`isl` uses the standard `autoconf` configure script. To run it, just type

```
./configure
```

optionally followed by some configure options. A complete list of options can be obtained by running

```
./configure --help
```

Below we discuss some of the more common options.

`isl` can optionally use `piplib`, but no `piplib` functionality is currently used by default. The `--with-piplib` option can be used to specify which `piplib` library to use, either an installed version (`system`), an externally built version (`build`) or no version (`no`). The option `build` is mostly useful in `configure` scripts of larger projects that bundle both `isl` and `piplib`.

--prefix

Installation prefix for `isl`

--with-gmp-prefix

Installation prefix for GMP (architecture-independent files).

--with-gmp-exec-prefix

Installation prefix for GMP (architecture-dependent files).

--with-piplib

Which copy of `piplib` to use, either `no` (default), `system` or `build`.

--with-piplib-prefix

Installation prefix for `system` `piplib` (architecture-independent files).

--with-piplib-exec-prefix

Installation prefix for `system` `piplib` (architecture-dependent files).

--with-piplib-builddir

Location where `build` `piplib` was built.

3. Compile

```
make
```

4. Install (optional)

```
make install
```

1.3 Library

1.3.1 Initialization

All manipulations of integer sets and relations occur within the context of an `isl_ctx`. A given `isl_ctx` can only be used within a single thread. All arguments of a function are required to have been allocated within the same context. There are currently no

functions available for moving an object from one `isl_ctx` to another `isl_ctx`. This means that there is currently no way of safely moving an object from one thread to another, unless the whole `isl_ctx` is moved.

An `isl_ctx` can be allocated using `isl_ctx_alloc` and freed using `isl_ctx_free`. All objects allocated within an `isl_ctx` should be freed before the `isl_ctx` itself is freed.

```
isl_ctx *isl_ctx_alloc();
void isl_ctx_free(isl_ctx *ctx);
```

1.3.2 Integers

All operations on integers, mainly the coefficients of the constraints describing the sets and relations, are performed in exact integer arithmetic using `GMP`. However, to allow future versions of `isl` to optionally support fixed integer arithmetic, all calls to `GMP` are wrapped inside `isl` specific macros. The basic type is `isl_int` and the following operations are available on this type. The meanings of these operations are essentially the same as their `GMP mpz` counterparts. As always with `GMP` types, `isl_ints` need to be initialized with `isl_int_init` before they can be used and they need to be released with `isl_int_clear` after the last use.

`isl_int_init(i)`

`isl_int_clear(i)`

`isl_int_set(r,i)`

`isl_int_set_si(r,i)`

`isl_int_abs(r,i)`

`isl_int_neg(r,i)`

`isl_int_swap(i,j)`

`isl_int_swap_or_set(i,j)`

`isl_int_add_ui(r,i,j)`

`isl_int_sub_ui(r,i,j)`

`isl_int_add(r,i,j)`

`isl_int_sub(r,i,j)`

`isl_int_mul(r,i,j)`

`isl_int_mul_ui(r,i,j)`

`isl_int_addmul(r,i,j)`

`isl_int_submul(r,i,j)`

isl_int_gcd(r,i,j)
isl_int_lcm(r,i,j)
isl_int_divexact(r,i,j)
isl_int_cdiv_q(r,i,j)
isl_int_fdiv_q(r,i,j)
isl_int_fdiv_r(r,i,j)
isl_int_fdiv_q_ui(r,i,j)
isl_int_read(r,s)
isl_int_print(out,i,width)
isl_int_sgn(i)
isl_int_cmp(i,j)
isl_int_cmp_si(i,si)
isl_int_eq(i,j)
isl_int_ne(i,j)
isl_int_lt(i,j)
isl_int_le(i,j)
isl_int_gt(i,j)
isl_int_ge(i,j)
isl_int_abs_eq(i,j)
isl_int_abs_ne(i,j)
isl_int_abs_lt(i,j)
isl_int_abs_gt(i,j)
isl_int_abs_ge(i,j)
isl_int_is_zero(i)
isl_int_is_one(i)
isl_int_is_negone(i)
isl_int_is_pos(i)
isl_int_is_neg(i)

isl_int_is_nonpos(i)

isl_int_is_nonneg(i)

isl_int_is_divisible_by(i,j)

1.3.3 Sets and Relations

isl uses four types of objects for representing sets and relations, `isl_basic_set`, `isl_basic_map`, `isl_set` and `isl_map`. `isl_basic_set` and `isl_basic_map` represent sets and relations that can be described as a conjunction of affine constraints, while `isl_set` and `isl_map` represent unions of `isl_basic_sets` and `isl_basic_maps`, respectively. The difference between sets and relations (maps) is that sets have one set of variables, while relations have two sets of variables, input variables and output variables.

1.3.4 Memory Management

Since a high-level operation on sets and/or relations usually involves several substeps and since the user is usually not interested in the intermediate results, most functions that return a new object will also release all the objects passed as arguments. If the user still wants to use one or more of these arguments after the function call, she should pass along a copy of the object rather than the object itself. The user is then responsible for make sure that the original object gets used somewhere else or is explicitly freed.

The arguments and return values of all documents functions are annotated to make clear which arguments are released and which arguments are preserved. In particular, the following annotations are used

`__isl_give`

`__isl_give` means that a new object is returned. The user should make sure that the returned pointer is used exactly once as a value for an `__isl_take` argument. In between, it can be used as a value for as many `__isl_keep` arguments as the user likes. There is one exception, and that is the case where the pointer returned is NULL. In this case, the user is free to use it as an `__isl_take` argument or not.

`__isl_take`

`__isl_take` means that the object the argument points to is taken over by the function and may no longer be used by the user as an argument to any other function. The pointer value must be one returned by a function returning an `__isl_give` pointer. If the user passes in a NULL value, then this will be treated as an error in the sense that the function will not perform its usual operation. However, it will still make sure that all the the other `__isl_take` arguments are released.

`__isl_keep`

`__isl_keep` means that the function will only use the object temporarily. After the function has finished, the user can still use it as an argument to other functions. A NULL value will be treated in the same way as a NULL value for an `__isl_take` argument.

1.3.5 Dimension Specifications

Whenever a new set or relation is created from scratch, its dimension needs to be specified using an `isl_dim`.

```
#include <isl_dim.h>
__isl_give isl_dim *isl_dim_alloc(isl_ctx *ctx,
    unsigned nparam, unsigned n_in, unsigned n_out);
__isl_give isl_dim *isl_dim_set_alloc(isl_ctx *ctx,
    unsigned nparam, unsigned dim);
__isl_give isl_dim *isl_dim_copy(__isl_keep isl_dim *dim);
void isl_dim_free(__isl_take isl_dim *dim);
unsigned isl_dim_size(__isl_keep isl_dim *dim,
    enum isl_dim_type type);
```

The dimension specification used for creating a set needs to be created using `isl_dim_set_alloc`, while that for creating a relation needs to be created using `isl_dim_alloc`. `isl_dim_size` can be used to find out the number of dimensions of each type in a dimension specification, where `type` may be `isl_dim_param`, `isl_dim_in` (only for relations), `isl_dim_out` (only for relations), `isl_dim_set` (only for sets) or `isl_dim_all`.

It is often useful to create objects that live in the same space as some other object. This can be accomplished by creating the new objects (see [Creating New Sets and Relations](#) or [Creating New \(Piecewise\) Quasipolynomials](#)) based on the dimension specification of the original object.

```
#include <isl_set.h>
__isl_give isl_dim *isl_basic_set_get_dim(
    __isl_keep isl_basic_set *bset);
__isl_give isl_dim *isl_set_get_dim(__isl_keep isl_set *set);

#include <isl_map.h>
__isl_give isl_dim *isl_basic_map_get_dim(
    __isl_keep isl_basic_map *bmap);
__isl_give isl_dim *isl_map_get_dim(__isl_keep isl_map *map);

#include <isl_polynomial.h>
__isl_give isl_dim *isl_qpolynomial_get_dim(
    __isl_keep isl_qpolynomial *qp);
__isl_give isl_dim *isl_pw_qpolynomial_get_dim(
    __isl_keep isl_pw_qpolynomial *pwqp);
```

The names of the individual dimensions may be set or read off using the following functions.

```

#include <isl_dim.h>
__isl_give isl_dim *isl_dim_set_name(__isl_take isl_dim *dim,
                                     enum isl_dim_type type, unsigned pos,
                                     __isl_keep const char *name);
__isl_keep const char *isl_dim_get_name(__isl_keep isl_dim *dim,
                                         enum isl_dim_type type, unsigned pos);

```

Note that `isl_dim_get_name` returns a pointer to some internal data structure, so the result can only be used while the corresponding `isl_dim` is alive. Also note that every function that operates on two sets or relations requires that both arguments have the same parameters. This also means that if one of the arguments has named parameters, then the other needs to have named parameters too and the names need to match.

1.3.6 Input and Output

`isl` supports its own input/output format, which is similar to the `Omega` format, but also supports the `PolyLib` format in some cases.

isl format

The `isl` format is similar to that of `Omega`, but has a different syntax for describing the parameters and allows for the definition of an existentially quantified variable as the integer division of an affine expression. For example, the set of integers `i` between `0` and `n` such that `i % 10 <= 6` can be described as

$$[n] \rightarrow \{ [i] : \text{exists } (a = [i/10] : 0 \leq i \text{ and } i \leq n \text{ and } i - 10a \leq 6) \}$$

A set or relation can have several disjuncts, separated by the keyword `or`. Each disjunct is either a conjunction of constraints or a projection (`exists`) of a conjunction of constraints. The constraints are separated by the keyword `and`.

PolyLib format

If the represented set is a union, then the first line contains a single number representing the number of disjuncts. Otherwise, a line containing the number 1 is optional.

Each disjunct is represented by a matrix of constraints. The first line contains two numbers representing the number of rows and columns, where the number of rows is equal to the number of constraints and the number of columns is equal to two plus the number of variables. The following lines contain the actual rows of the constraint matrix. In each row, the first column indicates whether the constraint is an equality (0) or inequality (1). The final column corresponds to the constant term.

If the set is parametric, then the coefficients of the parameters appear in the last columns before the constant column. The coefficients of any existentially quantified variables appear between those of the set variables and those of the parameters.

Input

```
#include <isl_set.h>
__isl_give isl_basic_set *isl_basic_set_read_from_file(
    isl_ctx *ctx, FILE *input, int nparam);
__isl_give isl_basic_set *isl_basic_set_read_from_str(
    isl_ctx *ctx, const char *str, int nparam);
__isl_give isl_set *isl_set_read_from_file(isl_ctx *ctx,
    FILE *input, int nparam);
__isl_give isl_set *isl_set_read_from_str(isl_ctx *ctx,
    const char *str, int nparam);

#include <isl_map.h>
__isl_give isl_basic_map *isl_basic_map_read_from_file(
    isl_ctx *ctx, FILE *input, int nparam);
__isl_give isl_basic_map *isl_basic_map_read_from_str(
    isl_ctx *ctx, const char *str, int nparam);
__isl_give isl_map *isl_map_read_from_file(
    struct isl_ctx *ctx, FILE *input, int nparam);
__isl_give isl_map *isl_map_read_from_str(isl_ctx *ctx,
    const char *str, int nparam);
```

The input format is autodetected and may be either the PolyLib format or the isl format. `nparam` specifies how many of the final columns in the PolyLib format correspond to parameters. If input is given in the isl format, then the number of parameters needs to be equal to `nparam`. If `nparam` is negative, then any number of parameters is accepted in the isl format and zero parameters are assumed in the PolyLib format.

Output

Before anything can be printed, an `isl_printer` needs to be created.

```
__isl_give isl_printer *isl_printer_to_file(isl_ctx *ctx,
    FILE *file);
__isl_give isl_printer *isl_printer_to_str(isl_ctx *ctx);
void isl_printer_free(__isl_take isl_printer *printer);
__isl_give char *isl_printer_get_str(
    __isl_keep isl_printer *printer);
```

The behavior of the printer can be modified in various ways

```
__isl_give isl_printer *isl_printer_set_output_format(
    __isl_take isl_printer *p, int output_format);
__isl_give isl_printer *isl_printer_set_indent(
    __isl_take isl_printer *p, int indent);
__isl_give isl_printer *isl_printer_set_prefix(
    __isl_take isl_printer *p, const char *prefix);
```

```

__isl_give isl_printer *isl_printer_set_suffix(
    __isl_take isl_printer *p, const char *suffix);

```

The `output_format` may be either `ISL_FORMAT_ISL`, `ISL_FORMAT_OMEGA` or `ISL_FORMAT_POLYLIB` and defaults to `ISL_FORMAT_ISL`. Each line in the output is indented by `indent` spaces (default: 0), prefixed by `prefix` and suffixed by `suffix`. In the PolyLib format output, the coefficients of the existentially quantified variables appear between those of the set variables and those of the parameters.

To actually print something, use

```

#include <isl_set.h>
__isl_give isl_printer *isl_printer_print_basic_set(
    __isl_take isl_printer *printer,
    __isl_keep isl_basic_set *bset);
__isl_give isl_printer *isl_printer_print_set(
    __isl_take isl_printer *printer,
    __isl_keep isl_set *set);

#include <isl_map.h>
__isl_give isl_printer *isl_printer_print_basic_map(
    __isl_take isl_printer *printer,
    __isl_keep isl_basic_map *bmap);
__isl_give isl_printer *isl_printer_print_map(
    __isl_take isl_printer *printer,
    __isl_keep isl_map *map);

```

When called on a file printer, the following function flushes the file. When called on a string printer, the buffer is cleared.

```

__isl_give isl_printer *isl_printer_flush(
    __isl_take isl_printer *p);

```

1.3.7 Creating New Sets and Relations

`isl` has functions for creating some standard sets and relations.

- Empty sets and relations

```

__isl_give isl_basic_set *isl_basic_set_empty(
    __isl_take isl_dim *dim);
__isl_give isl_basic_map *isl_basic_map_empty(
    __isl_take isl_dim *dim);
__isl_give isl_set *isl_set_empty(
    __isl_take isl_dim *dim);
__isl_give isl_map *isl_map_empty(
    __isl_take isl_dim *dim);

```

- Universe sets and relations

```

__isl_give isl_basic_set *isl_basic_set_universe(
    __isl_take isl_dim *dim);
__isl_give isl_basic_map *isl_basic_map_universe(
    __isl_take isl_dim *dim);
__isl_give isl_set *isl_set_universe(
    __isl_take isl_dim *dim);
__isl_give isl_map *isl_map_universe(
    __isl_take isl_dim *dim);

```

- Identity relations

```

__isl_give isl_basic_map *isl_basic_map_identity(
    __isl_take isl_dim *set_dim);
__isl_give isl_map *isl_map_identity(
    __isl_take isl_dim *set_dim);

```

These functions take a dimension specification for a **set** and return an identity relation between two such sets.

- Lexicographic order

```

__isl_give isl_map *isl_map_lex_lt(
    __isl_take isl_dim *set_dim);
__isl_give isl_map *isl_map_lex_le(
    __isl_take isl_dim *set_dim);
__isl_give isl_map *isl_map_lex_gt(
    __isl_take isl_dim *set_dim);
__isl_give isl_map *isl_map_lex_ge(
    __isl_take isl_dim *set_dim);
__isl_give isl_map *isl_map_lex_lt_first(
    __isl_take isl_dim *dim, unsigned n);
__isl_give isl_map *isl_map_lex_le_first(
    __isl_take isl_dim *dim, unsigned n);
__isl_give isl_map *isl_map_lex_gt_first(
    __isl_take isl_dim *dim, unsigned n);
__isl_give isl_map *isl_map_lex_ge_first(
    __isl_take isl_dim *dim, unsigned n);

```

The first four functions take a dimension specification for a **set** and return relations that express that the elements in the domain are lexicographically less (`isl_map_lex_lt`), less or equal (`isl_map_lex_le`), greater (`isl_map_lex_gt`) or greater or equal (`isl_map_lex_ge`) than the elements in the range. The last four functions take a dimension specification for a map and return relations that express that the first `n` dimensions in the domain are lexicographically less (`isl_map_lex_lt_first`), less or equal (`isl_map_lex_le_first`), greater (`isl_map_lex_gt_first`) or greater or equal (`isl_map_lex_ge_first`) than the first `n` dimensions in the range.

A basic set or relation can be converted to a set or relation using the following functions.

```
__isl_give isl_set *isl_set_from_basic_set(
    __isl_take isl_basic_set *bset);
__isl_give isl_map *isl_map_from_basic_map(
    __isl_take isl_basic_map *bmap);
```

Sets and relations can be copied and freed again using the following functions.

```
__isl_give isl_basic_set *isl_basic_set_copy(
    __isl_keep isl_basic_set *bset);
__isl_give isl_set *isl_set_copy(__isl_keep isl_set *set);
__isl_give isl_basic_map *isl_basic_map_copy(
    __isl_keep isl_basic_map *bmap);
__isl_give isl_map *isl_map_copy(__isl_keep isl_map *map);
void isl_basic_set_free(__isl_take isl_basic_set *bset);
void isl_set_free(__isl_take isl_set *set);
void isl_basic_map_free(__isl_take isl_basic_map *bmap);
void isl_map_free(__isl_take isl_map *map);
```

Other sets and relations can be constructed by starting from a universe set or relation, adding equality and/or inequality constraints and then projecting out the existentially quantified variables, if any. Constraints can be constructed, manipulated and added to basic sets and relations using the following functions.

```
#include <isl_constraint.h>
__isl_give isl_constraint *isl_equality_alloc(
    __isl_take isl_dim *dim);
__isl_give isl_constraint *isl_inequality_alloc(
    __isl_take isl_dim *dim);
void isl_constraint_set_constant(
    __isl_keep isl_constraint *constraint, isl_int v);
void isl_constraint_set_coefficient(
    __isl_keep isl_constraint *constraint,
    enum isl_dim_type type, int pos, isl_int v);
__isl_give isl_basic_map *isl_basic_map_add_constraint(
    __isl_take isl_basic_map *bmap,
    __isl_take isl_constraint *constraint);
__isl_give isl_basic_set *isl_basic_set_add_constraint(
    __isl_take isl_basic_set *bset,
    __isl_take isl_constraint *constraint);
```

For example, to create a set containing the even integers between 10 and 42, you would use the following code.

```
isl_int v;
```

```

struct isl_dim *dim;
struct isl_constraint *c;
struct isl_basic_set *bset;

isl_int_init(v);
dim = isl_dim_set_alloc(ctx, 0, 2);
bset = isl_basic_set_universe(isl_dim_copy(dim));

c = isl_equality_alloc(isl_dim_copy(dim));
isl_int_set_si(v, -1);
isl_constraint_set_coefficient(c, isl_dim_set, 0, v);
isl_int_set_si(v, 2);
isl_constraint_set_coefficient(c, isl_dim_set, 1, v);
bset = isl_basic_set_add_constraint(bset, c);

c = isl_inequality_alloc(isl_dim_copy(dim));
isl_int_set_si(v, -10);
isl_constraint_set_constant(c, v);
isl_int_set_si(v, 1);
isl_constraint_set_coefficient(c, isl_dim_set, 0, v);
bset = isl_basic_set_add_constraint(bset, c);

c = isl_inequality_alloc(dim);
isl_int_set_si(v, 42);
isl_constraint_set_constant(c, v);
isl_int_set_si(v, -1);
isl_constraint_set_coefficient(c, isl_dim_set, 0, v);
bset = isl_basic_set_add_constraint(bset, c);

bset = isl_basic_set_project_out(bset, isl_dim_set, 1, 1);

isl_int_clear(v);

```

Or, alternatively,

```

struct isl_basic_set *bset;
bset = isl_basic_set_read_from_str(ctx,
    "{[i] : exists (a : i = 2a and i >= 10 and i <= 42)}", -1);

```

1.3.8 Inspecting Sets and Relations

Usually, the user should not have to care about the actual constraints of the sets and maps, but should instead apply the abstract operations explained in the following sections. Occasionally, however, it may be required to inspect the individual coefficients of the constraints. This section explains how to do so. In these cases, it may also be useful to have `isl` compute an explicit representation of the existentially quantified variables.

```

__isl_give isl_set *isl_set_compute_divs(
    __isl_take isl_set *set);
__isl_give isl_map *isl_map_compute_divs(
    __isl_take isl_map *map);

```

This explicit representation defines the existentially quantified variables as integer divisions of the other variables, possibly including earlier existentially quantified variables. An explicitly represented existentially quantified variable therefore has a unique value when the values of the other variables are known. If, furthermore, the same existentials, i.e., existentials with the same explicit representations, should appear in the same order in each of the disjuncts of a set or map, then the user should call either of the following functions.

```

__isl_give isl_set *isl_set_align_divs(
    __isl_take isl_set *set);
__isl_give isl_map *isl_map_align_divs(
    __isl_take isl_map *map);

```

To iterate over all the basic sets or maps in a set or map, use

```

int isl_set_foreach_basic_set(__isl_keep isl_set *set,
    int (*fn)(__isl_take isl_basic_set *bset, void *user),
    void *user);
int isl_map_foreach_basic_map(__isl_keep isl_map *map,
    int (*fn)(__isl_take isl_basic_map *bmap, void *user),
    void *user);

```

The callback function `fn` should return 0 if successful and -1 if an error occurs. In the latter case, or if any other error occurs, the above functions will return -1.

It should be noted that `isl` does not guarantee that the basic sets or maps passed to `fn` are disjoint. If this is required, then the user should call one of the following functions first.

```

__isl_give isl_set *isl_set_make_disjoint(
    __isl_take isl_set *set);
__isl_give isl_map *isl_map_make_disjoint(
    __isl_take isl_map *map);

```

To iterate over the constraints of a basic set or map, use

```

#include <isl_constraint.h>

int isl_basic_map_foreach_constraint(
    __isl_keep isl_basic_map *bmap,
    int (*fn)(__isl_take isl_constraint *c, void *user),
    void *user);
void isl_constraint_free(struct isl_constraint *c);

```

Again, the callback function `fn` should return 0 if successful and -1 if an error occurs. In the latter case, or if any other error occurs, the above functions will return -1. The constraint `c` represents either an equality or an inequality. Use the following function to find out whether a constraint represents an equality. If not, it represents an inequality.

```
int isl_constraint_is_equality(
    __isl_keep isl_constraint *constraint);
```

The coefficients of the constraints can be inspected using the following functions.

```
void isl_constraint_get_constant(
    __isl_keep isl_constraint *constraint, isl_int *v);
void isl_constraint_get_coefficient(
    __isl_keep isl_constraint *constraint,
    enum isl_dim_type type, int pos, isl_int *v);
```

The explicit representations of the existentially quantified variables can be inspected using the following functions. Note that the user is only allowed to use these functions if the inspected set or map is the result of a call to `isl_set_compute_divs` or `isl_map_compute_divs`.

```
__isl_give isl_div *isl_constraint_div(
    __isl_keep isl_constraint *constraint, int pos);
void isl_div_get_constant(__isl_keep isl_div *div,
    isl_int *v);
void isl_div_get_denominator(__isl_keep isl_div *div,
    isl_int *v);
void isl_div_get_coefficient(__isl_keep isl_div *div,
    enum isl_dim_type type, int pos, isl_int *v);
```

1.3.9 Properties

Unary Properties

- Emptiness

The following functions test whether the given set or relation contains any integer points. The “fast” variants do not perform any computations, but simply check if the given set or relation is already known to be empty.

```
int isl_basic_set_fast_is_empty(__isl_keep isl_basic_set *bset);
int isl_basic_set_is_empty(__isl_keep isl_basic_set *bset);
int isl_set_is_empty(__isl_keep isl_set *set);
int isl_basic_map_fast_is_empty(__isl_keep isl_basic_map *bmap);
int isl_basic_map_is_empty(__isl_keep isl_basic_map *bmap);
int isl_map_fast_is_empty(__isl_keep isl_map *map);
int isl_map_is_empty(__isl_keep isl_map *map);
```

- Universality

```
int isl_basic_set_is_universe(__isl_keep isl_basic_set *bset);
int isl_basic_map_is_universe(__isl_keep isl_basic_map *bmap);
int isl_set_fast_is_universe(__isl_keep isl_set *set);
```

- Single-valuedness

```
int isl_map_is_single_valued(__isl_keep isl_map *map);
```

- Bijectivity

```
int isl_map_is_bijective(__isl_keep isl_map *map);
```

Binary Properties

- Equality

```
int isl_set_fast_is_equal(__isl_keep isl_set *set1,
    __isl_keep isl_set *set2);
int isl_set_is_equal(__isl_keep isl_set *set1,
    __isl_keep isl_set *set2);
int isl_map_is_equal(__isl_keep isl_map *map1,
    __isl_keep isl_map *map2);
int isl_map_fast_is_equal(__isl_keep isl_map *map1,
    __isl_keep isl_map *map2);
int isl_basic_map_is_equal(
    __isl_keep isl_basic_map *bmap1,
    __isl_keep isl_basic_map *bmap2);
```

- Disjointness

```
int isl_set_fast_is_disjoint(__isl_keep isl_set *set1,
    __isl_keep isl_set *set2);
```

- Subset

```
int isl_set_is_subset(__isl_keep isl_set *set1,
    __isl_keep isl_set *set2);
int isl_set_is_strict_subset(
    __isl_keep isl_set *set1,
    __isl_keep isl_set *set2);
int isl_basic_map_is_subset(
    __isl_keep isl_basic_map *bmap1,
    __isl_keep isl_basic_map *bmap2);
int isl_basic_map_is_strict_subset(
```

```

        __isl_keep isl_basic_map *bmap1,
        __isl_keep isl_basic_map *bmap2);
int isl_map_is_subset(
    __isl_keep isl_map *map1,
    __isl_keep isl_map *map2);
int isl_map_is_strict_subset(
    __isl_keep isl_map *map1,
    __isl_keep isl_map *map2);

```

1.3.10 Unary Operations

- Complement

```

__isl_give isl_set *isl_set_complement(
    __isl_take isl_set *set);

```

- Inverse map

```

__isl_give isl_basic_map *isl_basic_map_reverse(
    __isl_take isl_basic_map *bmap);
__isl_give isl_map *isl_map_reverse(
    __isl_take isl_map *map);

```

- Projection

```

__isl_give isl_basic_set *isl_basic_set_project_out(
    __isl_take isl_basic_set *bset,
    enum isl_dim_type type, unsigned first, unsigned n);
__isl_give isl_basic_map *isl_basic_map_project_out(
    __isl_take isl_basic_map *bmap,
    enum isl_dim_type type, unsigned first, unsigned n);
__isl_give isl_set *isl_set_project_out(__isl_take isl_set *set,
    enum isl_dim_type type, unsigned first, unsigned n);
__isl_give isl_map *isl_map_project_out(__isl_take isl_map *map,
    enum isl_dim_type type, unsigned first, unsigned n);
__isl_give isl_basic_set *isl_basic_map_domain(
    __isl_take isl_basic_map *bmap);
__isl_give isl_basic_set *isl_basic_map_range(
    __isl_take isl_basic_map *bmap);
__isl_give isl_set *isl_map_domain(
    __isl_take isl_map *bmap);
__isl_give isl_set *isl_map_range(
    __isl_take isl_map *map);

```

- Deltas

```

__isl_give isl_basic_set *isl_basic_map_deltas(
    __isl_take isl_basic_map *bmap);
__isl_give isl_set *isl_map_deltas(__isl_take isl_map *map);

```

These functions return a (basic) set containing the differences between image elements and corresponding domain elements in the input.

- Coalescing

Simplify the representation of a set or relation by trying to combine pairs of basic sets or relations into a single basic set or relation.

```

__isl_give isl_set *isl_set_coalesce(__isl_take isl_set *set);
__isl_give isl_map *isl_map_coalesce(__isl_take isl_map *map);

```

- Convex hull

```

__isl_give isl_basic_set *isl_set_convex_hull(
    __isl_take isl_set *set);
__isl_give isl_basic_map *isl_map_convex_hull(
    __isl_take isl_map *map);

```

If the input set or relation has any existentially quantified variables, then the result of these operations is currently undefined.

- Simple hull

```

__isl_give isl_basic_set *isl_set_simple_hull(
    __isl_take isl_set *set);
__isl_give isl_basic_map *isl_map_simple_hull(
    __isl_take isl_map *map);

```

These functions compute a single basic set or relation that contains the whole input set or relation. In particular, the output is described by translates of the constraints describing the basic sets or relations in the input.

(See Section 2.2.)

- Affine hull

```

__isl_give isl_basic_set *isl_basic_set_affine_hull(
    __isl_take isl_basic_set *bset);
__isl_give isl_basic_set *isl_set_affine_hull(
    __isl_take isl_set *set);
__isl_give isl_basic_map *isl_basic_map_affine_hull(
    __isl_take isl_basic_map *bmap);
__isl_give isl_basic_map *isl_map_affine_hull(
    __isl_take isl_map *map);

```

- Power

```
__isl_give isl_map *isl_map_power(__isl_take isl_map *map,
    unsigned param, int *exact);
```

Compute a parametric representation for all positive powers k of `map`. The power k is equated to the parameter at position `param`. The result may be an overapproximation. If the result is exact, then `*exact` is set to 1. The current implementation only produces exact results for particular cases of piecewise translations (i.e., piecewise uniform dependences).

- Transitive closure

```
__isl_give isl_map *isl_map_transitive_closure(
    __isl_take isl_map *map, int *exact);
```

Compute the transitive closure of `map`. The result may be an overapproximation. If the result is known to be exact, then `*exact` is set to 1. The current implementation only produces exact results for particular cases of piecewise translations (i.e., piecewise uniform dependences).

- Reaching path lengths

```
__isl_give isl_map *isl_map_reaching_path_lengths(
    __isl_take isl_map *map, int *exact);
```

Compute a relation that maps each element in the range of `map` to the lengths of all paths composed of edges in `map` that end up in the given element. The result may be an overapproximation. If the result is known to be exact, then `*exact` is set to 1. To compute the *maximal* path length, the resulting relation should be postprocessed by `isl_map_lexmax`. In particular, if the input relation is a dependence relation (mapping sources to sinks), then the maximal path length corresponds to the free schedule. Note, however, that `isl_map_lexmax` expects the maximum to be finite, so if the path lengths are unbounded (possibly due to the overapproximation), then you will get an error message.

1.3.11 Binary Operations

The two arguments of a binary operation not only need to live in the same `isl_ctx`, they currently also need to have the same (number of) parameters.

Basic Operations

- Intersection

```

__isl_give isl_basic_set *isl_basic_set_intersect(
    __isl_take isl_basic_set *bset1,
    __isl_take isl_basic_set *bset2);
__isl_give isl_set *isl_set_intersect(
    __isl_take isl_set *set1,
    __isl_take isl_set *set2);
__isl_give isl_basic_map *isl_basic_map_intersect_domain(
    __isl_take isl_basic_map *bmap,
    __isl_take isl_basic_set *bset);
__isl_give isl_basic_map *isl_basic_map_intersect_range(
    __isl_take isl_basic_map *bmap,
    __isl_take isl_basic_set *bset);
__isl_give isl_basic_map *isl_basic_map_intersect(
    __isl_take isl_basic_map *bmap1,
    __isl_take isl_basic_map *bmap2);
__isl_give isl_map *isl_map_intersect_domain(
    __isl_take isl_map *map,
    __isl_take isl_set *set);
__isl_give isl_map *isl_map_intersect_range(
    __isl_take isl_map *map,
    __isl_take isl_set *set);
__isl_give isl_map *isl_map_intersect(
    __isl_take isl_map *map1,
    __isl_take isl_map *map2);

```

- Union

```

__isl_give isl_set *isl_basic_set_union(
    __isl_take isl_basic_set *bset1,
    __isl_take isl_basic_set *bset2);
__isl_give isl_map *isl_basic_map_union(
    __isl_take isl_basic_map *bmap1,
    __isl_take isl_basic_map *bmap2);
__isl_give isl_set *isl_set_union(
    __isl_take isl_set *set1,
    __isl_take isl_set *set2);
__isl_give isl_map *isl_map_union(
    __isl_take isl_map *map1,
    __isl_take isl_map *map2);

```

- Set difference

```

__isl_give isl_set *isl_set_subtract(
    __isl_take isl_set *set1,
    __isl_take isl_set *set2);
__isl_give isl_map *isl_map_subtract(

```

```

__isl_take isl_map *map1,
__isl_take isl_map *map2);

```

- Application

```

__isl_give isl_basic_set *isl_basic_set_apply(
    __isl_take isl_basic_set *bset,
    __isl_take isl_basic_map *bmap);
__isl_give isl_set *isl_set_apply(
    __isl_take isl_set *set,
    __isl_take isl_map *map);
__isl_give isl_basic_map *isl_basic_map_apply_domain(
    __isl_take isl_basic_map *bmap1,
    __isl_take isl_basic_map *bmap2);
__isl_give isl_basic_map *isl_basic_map_apply_range(
    __isl_take isl_basic_map *bmap1,
    __isl_take isl_basic_map *bmap2);
__isl_give isl_map *isl_map_apply_domain(
    __isl_take isl_map *map1,
    __isl_take isl_map *map2);
__isl_give isl_map *isl_map_apply_range(
    __isl_take isl_map *map1,
    __isl_take isl_map *map2);

```

- Simplification

```

__isl_give isl_basic_set *isl_basic_set_gist(
    __isl_take isl_basic_set *bset,
    __isl_take isl_basic_set *context);
__isl_give isl_set *isl_set_gist(__isl_take isl_set *set,
    __isl_take isl_set *context);
__isl_give isl_basic_map *isl_basic_map_gist(
    __isl_take isl_basic_map *bmap,
    __isl_take isl_basic_map *context);
__isl_give isl_map *isl_map_gist(__isl_take isl_map *map,
    __isl_take isl_map *context);

```

The gist operation returns a set or relation that has the same intersection with the context as the input set or relation. Any implicit equality in the intersection is made explicit in the result, while all inequalities that are redundant with respect to the intersection are removed.

Lexicographic Optimization

Given a (basic) set `set` (or `bset`) and a zero-dimensional domain `dom`, the following functions compute a set that contains the lexicographic minimum or maximum of the

elements in `set` (or `bset`) for those values of the parameters that satisfy `dom`. If `empty` is not `NULL`, then `*empty` is assigned a set that contains the parameter values in `dom` for which `set` (or `bset`) has no elements. In other words, the union of the parameter values for which the result is non-empty and of `*empty` is equal to `dom`.

```

__isl_give isl_set *isl_basic_set_partial_lexmin(
    __isl_take isl_basic_set *bset,
    __isl_take isl_basic_set *dom,
    __isl_give isl_set **empty);
__isl_give isl_set *isl_basic_set_partial_lexmax(
    __isl_take isl_basic_set *bset,
    __isl_take isl_basic_set *dom,
    __isl_give isl_set **empty);
__isl_give isl_set *isl_set_partial_lexmin(
    __isl_take isl_set *set, __isl_take isl_set *dom,
    __isl_give isl_set **empty);
__isl_give isl_set *isl_set_partial_lexmax(
    __isl_take isl_set *set, __isl_take isl_set *dom,
    __isl_give isl_set **empty);

```

Given a (basic) set `set` (or `bset`), the following functions simply return a set containing the lexicographic minimum or maximum of the elements in `set` (or `bset`).

```

__isl_give isl_set *isl_basic_set_lexmin(
    __isl_take isl_basic_set *bset);
__isl_give isl_set *isl_basic_set_lexmax(
    __isl_take isl_basic_set *bset);
__isl_give isl_set *isl_set_lexmin(
    __isl_take isl_set *set);
__isl_give isl_set *isl_set_lexmax(
    __isl_take isl_set *set);

```

Given a (basic) relation map (or `bmap`) and a domain `dom`, the following functions compute a relation that maps each element of `dom` to the single lexicographic minimum or maximum of the elements that are associated to that same element in `map` (or `bmap`). If `empty` is not `NULL`, then `*empty` is assigned a set that contains the elements in `dom` that do not map to any elements in `map` (or `bmap`). In other words, the union of the domain of the result and of `*empty` is equal to `dom`.

```

__isl_give isl_map *isl_basic_map_partial_lexmax(
    __isl_take isl_basic_map *bmap,
    __isl_take isl_basic_set *dom,
    __isl_give isl_set **empty);
__isl_give isl_map *isl_basic_map_partial_lexmin(
    __isl_take isl_basic_map *bmap,
    __isl_take isl_basic_set *dom,
    __isl_give isl_set **empty);

```

```

__isl_give isl_map *isl_map_partial_lexmax(
    __isl_take isl_map *map, __isl_take isl_set *dom,
    __isl_give isl_set **empty);
__isl_give isl_map *isl_map_partial_lexmin(
    __isl_take isl_map *map, __isl_take isl_set *dom,
    __isl_give isl_set **empty);

```

Given a (basic) map map (or bmap), the following functions simply return a map mapping each element in the domain of map (or bmap) to the lexicographic minimum or maximum of all elements associated to that element.

```

__isl_give isl_map *isl_basic_map_lexmin(
    __isl_take isl_basic_map *bmap);
__isl_give isl_map *isl_basic_map_lexmax(
    __isl_take isl_basic_map *bmap);
__isl_give isl_map *isl_map_lexmin(
    __isl_take isl_map *map);
__isl_give isl_map *isl_map_lexmax(
    __isl_take isl_map *map);

```

1.3.12 Points

Points are elements of a set. They can be used to construct simple sets (boxes) or they can be used to represent the individual elements of a set. The zero point (the origin) can be created using

```

__isl_give isl_point *isl_point_zero(__isl_take isl_dim *dim);

```

The coordinates of a point can be inspected, set and changed using

```

void isl_point_get_coordinate(__isl_keep isl_point *pnt,
    enum isl_dim_type type, int pos, isl_int *v);
__isl_give isl_point *isl_point_set_coordinate(
    __isl_take isl_point *pnt,
    enum isl_dim_type type, int pos, isl_int v);

__isl_give isl_point *isl_point_add_ui(
    __isl_take isl_point *pnt,
    enum isl_dim_type type, int pos, unsigned val);
__isl_give isl_point *isl_point_sub_ui(
    __isl_take isl_point *pnt,
    enum isl_dim_type type, int pos, unsigned val);

```

Points can be copied or freed using

```

__isl_give isl_point *isl_point_copy(
    __isl_keep isl_point *pnt);
void isl_point_free(__isl_take isl_point *pnt);

```

A singleton set can be created from a point using

```
__isl_give isl_set *isl_set_from_point(
    __isl_take isl_point *pnt);
```

and a box can be created from two opposite extremal points using

```
__isl_give isl_set *isl_set_box_from_points(
    __isl_take isl_point *pnt1,
    __isl_take isl_point *pnt2);
```

All elements of a **bounded** set can be enumerated using the following function.

```
int isl_set_foreach_point(__isl_keep isl_set *set,
    int (*fn)(__isl_take isl_point *pnt, void *user),
    void *user);
```

The function `fn` is called for each integer point in `set` with as second argument the last argument of the `isl_set_foreach_point` call. The function `fn` should return `0` on success and `-1` on failure. In the latter case, `isl_set_foreach_point` will stop enumerating and return `-1` as well. If the enumeration is performed successfully and to completion, then `isl_set_foreach_point` returns `0`.

To obtain a single point of a set, use

```
__isl_give isl_point *isl_set_sample_point(
    __isl_take isl_set *set);
```

If `set` does not contain any (integer) points, then the resulting point will be “void”, a property that can be tested using

```
int isl_point_is_void(__isl_keep isl_point *pnt);
```

1.3.13 Piecewise Quasipolynomials

A piecewise quasipolynomial is a particular kind of function that maps a parametric point to a rational value. More specifically, a quasipolynomial is a polynomial expression in greatest integer parts of affine expressions of parameters and variables. A piecewise quasipolynomial is a subdivision of a given parametric domain into disjoint cells with a quasipolynomial associated to each cell. The value of the piecewise quasipolynomial at a given point is the value of the quasipolynomial associated to the cell that contains the point. Outside of the union of cells, the value is assumed to be zero. For example, the piecewise quasipolynomial

$$[n] \rightarrow \{ [x] \rightarrow ((1 + n) - x) : x \leq n \text{ and } x \geq 0 \}$$

maps x to $1 + n - x$ for values of x between 0 and n . Piecewise quasipolynomials are mainly used by the `barvinok` library for representing the number of elements in a parametric set or map. For example, the piecewise quasipolynomial above represents the number of points in the map

$$[n] \rightarrow \{ [x] \rightarrow [y] : x, y \geq 0 \text{ and } 0 \leq x + y \leq n \}$$

Printing (Piecewise) Quasipolynomials

Quasipolynomials and piecewise quasipolynomials can be printed using the following functions.

```
__isl_give isl_printer *isl_printer_print_qpolynomial(
    __isl_take isl_printer *p,
    __isl_keep isl_qpolynomial *qp);

__isl_give isl_printer *isl_printer_print_pw_qpolynomial(
    __isl_take isl_printer *p,
    __isl_keep isl_pw_qpolynomial *pwqp);
```

The output format of the printer needs to be set to either `ISL_FORMAT_ISL` or `ISL_FORMAT_C`.

Creating New (Piecewise) Quasipolynomials

Some simple quasipolynomials can be created using the following functions. More complicated quasipolynomials can be created by applying operations such as addition and multiplication on the resulting quasipolynomials

```
__isl_give isl_qpolynomial *isl_qpolynomial_zero(
    __isl_take isl_dim *dim);
__isl_give isl_qpolynomial *isl_qpolynomial_infty(
    __isl_take isl_dim *dim);
__isl_give isl_qpolynomial *isl_qpolynomial_neginfty(
    __isl_take isl_dim *dim);
__isl_give isl_qpolynomial *isl_qpolynomial_nan(
    __isl_take isl_dim *dim);
__isl_give isl_qpolynomial *isl_qpolynomial_rat_cst(
    __isl_take isl_dim *dim,
    const isl_int n, const isl_int d);
__isl_give isl_qpolynomial *isl_qpolynomial_div(
    __isl_take isl_div *div);
__isl_give isl_qpolynomial *isl_qpolynomial_var(
    __isl_take isl_dim *dim,
    enum isl_dim_type type, unsigned pos);
```

The zero piecewise quasipolynomial or a piecewise quasipolynomial with a single cell can be created using the following functions. Multiple of these single cell piecewise quasipolynomials can be combined to create more complicated piecewise quasipolynomials.

```
__isl_give isl_pw_qpolynomial *isl_pw_qpolynomial_zero(
    __isl_take isl_dim *dim);
__isl_give isl_pw_qpolynomial *isl_pw_qpolynomial_alloc(
    __isl_take isl_set *set,
    __isl_take isl_qpolynomial *qp);
```

Quasipolynomials can be copied and freed again using the following functions.

```
__isl_give isl_qpolynomial *isl_qpolynomial_copy(
    __isl_keep isl_qpolynomial *qp);
void isl_qpolynomial_free(__isl_take isl_qpolynomial *qp);

__isl_give isl_pw_qpolynomial *isl_pw_qpolynomial_copy(
    __isl_keep isl_pw_qpolynomial *pwqp);
void isl_pw_qpolynomial_free(
    __isl_take isl_pw_qpolynomial *pwqp);
```

Inspecting (Piecewise) Quasipolynomials

To iterate over the cells in a piecewise quasipolynomial, use either of the following two functions

```
int isl_pw_qpolynomial_foreach_piece(
    __isl_keep isl_pw_qpolynomial *pwqp,
    int (*fn)(__isl_take isl_set *set,
              __isl_take isl_qpolynomial *qp,
              void *user), void *user);
int isl_pw_qpolynomial_foreach_lifted_piece(
    __isl_keep isl_pw_qpolynomial *pwqp,
    int (*fn)(__isl_take isl_set *set,
              __isl_take isl_qpolynomial *qp,
              void *user), void *user);
```

As usual, the function `fn` should return `0` on success and `-1` on failure. The difference between `isl_pw_qpolynomial_foreach_piece` and `isl_pw_qpolynomial_foreach_lifted_piece` is that `isl_pw_qpolynomial_foreach_lifted_piece` will first compute unique representations for all existentially quantified variables and then turn these existentially quantified variables into extra set variables, adapting the associated quasipolynomial accordingly. This means that the set passed to `fn` will not have any existentially quantified variables, but that the dimensions of the sets may be different for different invocations of `fn`.

To iterate over all terms in a quasipolynomial, use

```
int isl_qpolynomial_foreach_term(
    __isl_keep isl_qpolynomial *qp,
    int (*fn)(__isl_take isl_term *term,
              void *user), void *user);
```

The terms themselves can be inspected and freed using these functions

```
unsigned isl_term_dim(__isl_keep isl_term *term,
    enum isl_dim_type type);
void isl_term_get_num(__isl_keep isl_term *term,
```

```

        isl_int *n);
void isl_term_get_den(__isl_keep isl_term *term,
                    isl_int *d);
int isl_term_get_exp(__isl_keep isl_term *term,
                    enum isl_dim_type type, unsigned pos);
__isl_give isl_div *isl_term_get_div(
    __isl_keep isl_term *term, unsigned pos);
void isl_term_free(__isl_take isl_term *term);

```

Each term is a product of parameters, set variables and integer divisions. The function `isl_term_get_exp` returns the exponent of a given dimensions in the given term. The `isl_int`s in the arguments of `isl_term_get_num` and `isl_term_get_den` need to have been initialized using `isl_int_init` before calling these functions.

Properties of (Piecewise) Quasipolynomials

To check whether a quasipolynomial is actually a constant, use the following function.

```

int isl_qupolynomial_is_cst(__isl_keep isl_qupolynomial *qp,
                          isl_int *n, isl_int *d);

```

If `qp` is a constant and if `n` and `d` are not `NULL` then the numerator and denominator of the constant are returned in `*n` and `*d`, respectively.

Operations on (Piecewise) Quasipolynomials

```

__isl_give isl_qupolynomial *isl_qupolynomial_neg(
    __isl_take isl_qupolynomial *qp);
__isl_give isl_qupolynomial *isl_qupolynomial_add(
    __isl_take isl_qupolynomial *qp1,
    __isl_take isl_qupolynomial *qp2);
__isl_give isl_qupolynomial *isl_qupolynomial_sub(
    __isl_take isl_qupolynomial *qp1,
    __isl_take isl_qupolynomial *qp2);
__isl_give isl_qupolynomial *isl_qupolynomial_mul(
    __isl_take isl_qupolynomial *qp1,
    __isl_take isl_qupolynomial *qp2);

__isl_give isl_pw_qupolynomial *isl_pw_qupolynomial_add(
    __isl_take isl_pw_qupolynomial *pwqp1,
    __isl_take isl_pw_qupolynomial *pwqp2);
__isl_give isl_pw_qupolynomial *isl_pw_qupolynomial_sub(
    __isl_take isl_pw_qupolynomial *pwqp1,
    __isl_take isl_pw_qupolynomial *pwqp2);
__isl_give isl_pw_qupolynomial *isl_pw_qupolynomial_add_disjoint(
    __isl_take isl_pw_qupolynomial *pwqp1,
    __isl_take isl_pw_qupolynomial *pwqp2);

```

```

__isl_give isl_pw_qpolynomial *isl_pw_qpolynomial_neg(
    __isl_take isl_pw_qpolynomial *pwqp);
__isl_give isl_pw_qpolynomial *isl_pw_qpolynomial_mul(
    __isl_take isl_pw_qpolynomial *pwqp1,
    __isl_take isl_pw_qpolynomial *pwqp2);

__isl_give isl_qpolynomial *isl_pw_qpolynomial_eval(
    __isl_take isl_pw_qpolynomial *pwqp,
    __isl_take isl_point *pnt);

__isl_give isl_set *isl_pw_qpolynomial_domain(
    __isl_take isl_pw_qpolynomial *pwqp);
__isl_give isl_pw_qpolynomial *isl_pw_qpolynomial_intersect_domain(
    __isl_take isl_pw_qpolynomial *pwqp,
    __isl_take isl_set *set);

__isl_give isl_pw_qpolynomial *isl_pw_qpolynomial_gist(
    __isl_take isl_pw_qpolynomial *pwqp,
    __isl_take isl_set *context);

```

The gist operation applies the gist operation to each of the cells in the domain of the input piecewise quasipolynomial. In future, the operation will also exploit the context to simplify the quasipolynomials associated to each cell.

1.3.14 Bounds on Piecewise Quasipolynomials and Piecewise Quasipolynomial Reductions

A piecewise quasipolynomial reduction is a piecewise reduction (or fold) of quasipolynomials. In particular, the reduction can be maximum or a minimum. The objects are mainly used to represent the result of an upper or lower bound on a quasipolynomial over its domain, i.e., as the result of the following function.

```

__isl_give isl_pw_qpolynomial_fold *isl_pw_qpolynomial_bound(
    __isl_take isl_pw_qpolynomial *pwqp,
    enum isl_fold type, int *tight);

```

The type argument may be either `isl_fold_min` or `isl_fold_max`. If `tight` is not `NULL`, then `*tight` is set to 1 if the returned bound is known to be tight, i.e., for each value of the parameters there is at least one element in the domain that reaches the bound.

A (piecewise) quasipolynomial reduction can be copied or freed using the following functions.

```

__isl_give isl_qpolynomial_fold *isl_qpolynomial_fold_copy(
    __isl_keep isl_qpolynomial_fold *fold);
__isl_give isl_pw_qpolynomial_fold *isl_pw_qpolynomial_fold_copy(
    __isl_keep isl_pw_qpolynomial_fold *pwf);

```

```

void isl_qpolynomial_fold_free(
    __isl_take isl_qpolynomial_fold *fold);
void isl_pw_qpolynomial_fold_free(
    __isl_take isl_pw_qpolynomial_fold *pwf);

```

Printing Piecewise Quasipolynomial Reductions

Piecewise quasipolynomial reductions can be printed using the following function.

```

__isl_give isl_printer *isl_printer_print_pw_qpolynomial_fold(
    __isl_take isl_printer *p,
    __isl_keep isl_pw_qpolynomial_fold *pwf);

```

The output format of the printer needs to be set to either `ISL_FORMAT_ISL` or `ISL_FORMAT_C`.

Inspecting (Piecewise) Quasipolynomial Reductions

To iterate over the cells in a piecewise quasipolynomial reduction, use either of the following two functions

```

int isl_pw_qpolynomial_fold_foreach_piece(
    __isl_keep isl_pw_qpolynomial_fold *pwf,
    int (*fn)(__isl_take isl_set *set,
              __isl_take isl_qpolynomial_fold *fold,
              void *user), void *user);
int isl_pw_qpolynomial_fold_foreach_lifted_piece(
    __isl_keep isl_pw_qpolynomial_fold *pwf,
    int (*fn)(__isl_take isl_set *set,
              __isl_take isl_qpolynomial_fold *fold,
              void *user), void *user);

```

See [Inspecting \(Piecewise\) Quasipolynomials](#) for an explanation of the difference between these two functions.

To iterate over all quasipolynomials in a reduction, use

```

int isl_qpolynomial_fold_foreach_qpolynomial(
    __isl_keep isl_qpolynomial_fold *fold,
    int (*fn)(__isl_take isl_qpolynomial *qp,
              void *user), void *user);

```

Operations on Piecewise Quasipolynomial Reductions

```

__isl_give isl_pw_qpolynomial_fold *isl_pw_qpolynomial_fold_add(
    __isl_take isl_pw_qpolynomial_fold *pwf1,
    __isl_take isl_pw_qpolynomial_fold *pwf2);

```

```

__isl_give isl_qpolynomial *isl_pw_qpolynomial_fold_eval(
    __isl_take isl_pw_qpolynomial_fold *pwf,
    __isl_take isl_point *pnt);

__isl_give isl_pw_qpolynomial_fold *isl_pw_qpolynomial_fold_coalesce(
    __isl_take isl_pw_qpolynomial_fold *pwf);

__isl_give isl_pw_qpolynomial_fold *isl_pw_qpolynomial_fold_gist(
    __isl_take isl_pw_qpolynomial_fold *pwf,
    __isl_take isl_set *context);

```

The gist operation applies the gist operation to each of the cells in the domain of the input piecewise quasipolynomial reduction. In future, the operation will also exploit the context to simplify the quasipolynomial reductions associated to each cell.

1.3.15 Dependence Analysis

isl contains specialized functionality for performing array dataflow analysis. That is, given a *sink* access relation and a collection of possible *source* access relations, isl can compute relations that describe for each iteration of the sink access, which iteration of which of the source access relations was the last to access the same data element before the given iteration of the sink access. To compute standard flow dependences, the sink should be a read, while the sources should be writes. If any of the source accesses are marked as being *may* accesses, then there will be a dependence to the last *must* access **and** to any *may* access that follows this last *must* access. In particular, if *all* sources are *may* accesses, then memory based dependence analysis is performed. If, on the other hand, all sources are *must* accesses, then value based dependence analysis is performed.

```

#include <isl_flow.h>

__isl_give isl_access_info *isl_access_info_alloc(
    __isl_take isl_map *sink,
    void *sink_user, isl_access_level_before fn,
    int max_source);
__isl_give isl_access_info *isl_access_info_add_source(
    __isl_take isl_access_info *acc,
    __isl_take isl_map *source, int must,
    void *source_user);

__isl_give isl_flow *isl_access_info_compute_flow(
    __isl_take isl_access_info *acc);

int isl_flow_foreach(__isl_keep isl_flow *deps,
    int (*fn)(__isl_take isl_map *dep, int must,
              void *dep_user, void *user),
    void *user);

```

```

__isl_give isl_set *isl_flow_get_no_source(
    __isl_keep isl_flow *deps, int must);
void isl_flow_free(__isl_take isl_flow *deps);

```

The function `isl_access_info_compute_flow` performs the actual dependence analysis. The other functions are used to construct the input for this function or to read off the output.

The input is collected in an `isl_access_info`, which can be created through a call to `isl_access_info_alloc`. The arguments to this functions are the sink access relation `sink`, a token `sink_user` used to identify the sink access to the user, a callback function for specifying the relative order of source and sink accesses, and the number of source access relations that will be added. The callback function has type `int (*)(void *first, void *second)`. The function is called with two user supplied tokens identifying either a source or the sink and it should return the shared nesting level and the relative order of the two accesses. In particular, let n be the number of loops shared by the two accesses. If `first` precedes `second` textually, then the function should return $2 * n + 1$; otherwise, it should return $2 * n$. The sources can be added to the `isl_access_info` by performing (at most) `max_source` calls to `isl_access_info_add_source`. `must` indicates whether the source is a *must* access or a *may* access. Note that a multi-valued access relation should only be marked *must* if every iteration in the domain of the relation accesses *all* elements in its image. The `source_user` token is again used to identify the source access. The range of the source access relation `source` should have the same dimension as the range of the sink access relation.

The result of the dependence analysis is collected in an `isl_flow`. There may be elements in the domain of the sink access for which no preceding source access could be found or for which all preceding sources are *may* accesses. The sets of these elements can be obtained through calls to `isl_flow_get_no_source`, the first with `must` set and the second with `must` unset. In the case of standard flow dependence analysis, with the sink a read and the sources *must* writes, the first set corresponds to the reads from uninitialized array elements and the second set is empty. The actual flow dependences can be extracted using `isl_flow_foreach`. This function will call the user-specified callback function `fn` for each **non-empty** dependence between a source and the sink. The callback function is called with four arguments, the actual flow dependence relation mapping source iterations to sink iterations, a boolean that indicates whether it is a *must* or *may* dependence, a token identifying the source and an additional `void *` with value equal to the third argument of the `isl_flow_foreach` call. A dependence is marked *must* if it originates from a *must* source and if it is not followed by any *may* sources.

After finishing with an `isl_flow`, the user should call `isl_flow_free` to free all associated memory.

1.3.16 Parametric Vertex Enumeration

The parametric vertex enumeration described in this section is mainly intended to be used internally and by the `barvinok` library.

```

#include <isl_vertices.h>
__isl_give isl_vertices *isl_basic_set_compute_vertices(
    __isl_keep isl_basic_set *bset);

```

The function `isl_basic_set_compute_vertices` performs the actual computation of the parametric vertices and the chamber decomposition and store the result in an `isl_vertices` object. This information can be queried by either iterating over all the vertices or iterating over all the chambers or cells and then iterating over all vertices that are active on the chamber.

```

int isl_vertices_foreach_vertex(
    __isl_keep isl_vertices *vertices,
    int (*fn)(__isl_take isl_vertex *vertex, void *user),
    void *user);

int isl_vertices_foreach_cell(
    __isl_keep isl_vertices *vertices,
    int (*fn)(__isl_take isl_cell *cell, void *user),
    void *user);

int isl_cell_foreach_vertex(__isl_keep isl_cell *cell,
    int (*fn)(__isl_take isl_vertex *vertex, void *user),
    void *user);

```

Other operations that can be performed on an `isl_vertices` object are the following.

```

isl_ctx *isl_vertices_get_ctx(
    __isl_keep isl_vertices *vertices);
int isl_vertices_get_n_vertices(
    __isl_keep isl_vertices *vertices);
void isl_vertices_free(__isl_take isl_vertices *vertices);

```

Vertices can be inspected and destroyed using the following functions.

```

isl_ctx *isl_vertex_get_ctx(__isl_keep isl_vertex *vertex);
int isl_vertex_get_id(__isl_keep isl_vertex *vertex);
__isl_give isl_basic_set *isl_vertex_get_domain(
    __isl_keep isl_vertex *vertex);
__isl_give isl_basic_set *isl_vertex_get_expr(
    __isl_keep isl_vertex *vertex);
void isl_vertex_free(__isl_take isl_vertex *vertex);

```

`isl_vertex_get_expr` returns a singleton parametric set describing the vertex, while `isl_vertex_get_domain` returns the activity domain of the vertex. Note that `isl_vertex_get_domain` and `isl_vertex_get_expr` return **rational** basic sets, so they should mainly be used for inspection and should not be mixed with integer sets.

Chambers can be inspected and destroyed using the following functions.

```

isl_ctx *isl_cell_get_ctx(__isl_keep isl_cell *cell);
__isl_give isl_basic_set *isl_cell_get_domain(
    __isl_keep isl_cell *cell);
void isl_cell_free(__isl_take isl_cell *cell);

```

1.4 Applications

Although `isl` is mainly meant to be used as a library, it also contains some basic applications that use some of the functionality of `isl`. The input may be specified in either the `isl` format or the PolyLib format.

1.4.1 `isl_polyhedron_sample`

`isl_polyhedron_sample` takes a polyhedron as input and prints an integer element of the polyhedron, if there is any. The first column in the output is the denominator and is always equal to 1. If the polyhedron contains no integer points, then a vector of length zero is printed.

1.4.2 `isl_pip`

`isl_pip` takes the same input as the `example` program from the `piplib` distribution, i.e., a set of constraints on the parameters, a line contains only -1 and finally a set of constraints on a parametric polyhedron. The coefficients of the parameters appear in the last columns (but before the final constant column). The output is the lexicographic minimum of the parametric polyhedron. As `isl` currently does not have its own output format, the output is just a dump of the internal state.

1.4.3 `isl_polyhedron_minimize`

`isl_polyhedron_minimize` computes the minimum of some linear or affine objective function over the integer points in a polyhedron. If an affine objective function is given, then the constant should appear in the last column.

1.4.4 `isl_polytope_scan`

Given a polytope, `isl_polytope_scan` prints all integer points in the polytope.

1.5 `isl-polylib`

The `isl-polylib` library provides the following functions for converting between `isl` objects and PolyLib objects. The library is distributed separately for licensing reasons.

```
#include <isl_set_polylib.h>
__isl_give isl_basic_set *isl_basic_set_new_from_polylib(
    Polyhedron *P, __isl_take isl_dim *dim);
Polyhedron *isl_basic_set_to_polylib(
    __isl_keep isl_basic_set *bset);
__isl_give isl_set *isl_set_new_from_polylib(Polyhedron *D,
    __isl_take isl_dim *dim);
Polyhedron *isl_set_to_polylib(__isl_keep isl_set *set);

#include <isl_map_polylib.h>
__isl_give isl_basic_map *isl_basic_map_new_from_polylib(
    Polyhedron *P, __isl_take isl_dim *dim);
__isl_give isl_map *isl_map_new_from_polylib(Polyhedron *D,
    __isl_take isl_dim *dim);
Polyhedron *isl_basic_map_to_polylib(
    __isl_keep isl_basic_map *bmap);
Polyhedron *isl_map_to_polylib(__isl_keep isl_map *map);
```

Chapter 2

Implementation Details

2.1 Sets and Relations

Definition 2.1.1 (Polyhedral Set) A polyhedral set S is a finite union of basic sets $S = \bigcup_i S_i$, each of which can be represented using affine constraints

$$S_i : \mathbb{Z}^n \rightarrow 2^{\mathbb{Z}^d} : \mathbf{s} \mapsto S_i(\mathbf{s}) = \{ \mathbf{x} \in \mathbb{Z}^d \mid \exists \mathbf{z} \in \mathbb{Z}^e : A\mathbf{x} + B\mathbf{s} + D\mathbf{z} + \mathbf{c} \geq \mathbf{0} \},$$

with $A \in \mathbb{Z}^{m \times d}$, $B \in \mathbb{Z}^{m \times n}$, $D \in \mathbb{Z}^{m \times e}$ and $\mathbf{c} \in \mathbb{Z}^m$.

Definition 2.1.2 (Parameter Domain of a Set) Let $S \in \mathbb{Z}^n \rightarrow 2^{\mathbb{Z}^d}$ be a set. The parameter domain of S is the set

$$\text{pdom } S := \{ \mathbf{s} \in \mathbb{Z}^n \mid S(\mathbf{s}) \neq \emptyset \}.$$

Definition 2.1.3 (Polyhedral Relation) A polyhedral relation R is a finite union of basic relations $R = \bigcup_i R_i$ of type $\mathbb{Z}^n \rightarrow 2^{\mathbb{Z}^{d_1+d_2}}$, each of which can be represented using affine constraints

$$R_i = \mathbf{s} \mapsto R_i(\mathbf{s}) = \{ \mathbf{x}_1 \rightarrow \mathbf{x}_2 \in \mathbb{Z}^{d_1} \times \mathbb{Z}^{d_2} \mid \exists \mathbf{z} \in \mathbb{Z}^e : A_1\mathbf{x}_1 + A_2\mathbf{x}_2 + B\mathbf{s} + D\mathbf{z} + \mathbf{c} \geq \mathbf{0} \},$$

with $A_i \in \mathbb{Z}^{m \times d_i}$, $B \in \mathbb{Z}^{m \times n}$, $D \in \mathbb{Z}^{m \times e}$ and $\mathbf{c} \in \mathbb{Z}^m$.

Definition 2.1.4 (Parameter Domain of a Relation) Let $R \in \mathbb{Z}^n \rightarrow 2^{\mathbb{Z}^{d+d}}$ be a relation. The parameter domain of R is the set

$$\text{pdom } R := \{ \mathbf{s} \in \mathbb{Z}^n \mid R(\mathbf{s}) \neq \emptyset \}.$$

Definition 2.1.5 (Domain of a Relation) Let $R \in \mathbb{Z}^n \rightarrow 2^{\mathbb{Z}^{d+d}}$ be a relation. The domain of R is the polyhedral set

$$\text{dom } R := \mathbf{s} \mapsto \{ \mathbf{x}_1 \in \mathbb{Z}^{d_1} \mid \exists \mathbf{x}_2 \in \mathbb{Z}^{d_2} : (\mathbf{x}_1, \mathbf{x}_2) \in R(\mathbf{s}) \}.$$

Definition 2.1.6 (Range of a Relation) Let $R \in \mathbb{Z}^n \rightarrow 2^{\mathbb{Z}^{d+d}}$ be a relation. The range of R is the polyhedral set

$$\text{ran } R := \mathbf{s} \mapsto \{ \mathbf{x}_2 \in \mathbb{Z}^{d_2} \mid \exists \mathbf{x}_1 \in \mathbb{Z}^{d_1} : (\mathbf{x}_1, \mathbf{x}_2) \in R(\mathbf{s}) \}.$$

Definition 2.1.7 (Composition of Relations) Let $R \in \mathbb{Z}^n \rightarrow 2^{\mathbb{Z}^{d_1+d_2}}$ and $S \in \mathbb{Z}^n \rightarrow 2^{\mathbb{Z}^{d_2+d_3}}$ be two relations, then the composition of R and S is defined as

$$S \circ R := \mathbf{s} \mapsto \{ \mathbf{x}_1 \rightarrow \mathbf{x}_3 \in \mathbb{Z}^{d_1} \times \mathbb{Z}^{d_3} \mid \exists \mathbf{x}_2 \in \mathbb{Z}^{d_2} : \mathbf{x}_1 \rightarrow \mathbf{x}_2 \in R(\mathbf{s}) \wedge \mathbf{x}_2 \rightarrow \mathbf{x}_3 \in S(\mathbf{s}) \}.$$

Definition 2.1.8 (Difference Set of a Relation) Let $R \in \mathbb{Z}^n \rightarrow 2^{\mathbb{Z}^{d+d}}$ be a relation. The difference set (ΔR) of R is the set of differences between image elements and the corresponding domain elements,

$$\Delta R := \mathbf{s} \mapsto \{ \boldsymbol{\delta} \in \mathbb{Z}^d \mid \exists \mathbf{x} \rightarrow \mathbf{y} \in R : \boldsymbol{\delta} = \mathbf{y} - \mathbf{x} \}$$

2.2 Simple Hull

It is sometimes useful to have a single basic set or basic relation that contains a given set or relation. For rational sets, the obvious choice would be to compute the (rational) convex hull. For integer sets, the obvious choice would be the integer hull. However, isl currently does not support an integer hull operation and even if it did, it would be fairly expensive to compute. The convex hull operation is supported, but it is also fairly expensive to compute given only an implicit representation.

Usually, it is not required to compute the exact integer hull, and an overapproximation of this hull is sufficient. The “simple hull” of a set is such an overapproximation and it is defined as the (inclusion-wise) smallest basic set that is described by constraints that are translates of the constraints in the input set. This means that the simple hull is relatively cheap to compute and that the number of constraints in the simple hull is no larger than the number of constraints in the input.

Definition 2.2.1 (Simple Hull of a Set) The simple hull of a set $S = \bigcup_{1 \leq i \leq v} S_i$, with

$$S : \mathbb{Z}^n \rightarrow 2^{\mathbb{Z}^d} : \mathbf{s} \mapsto S(\mathbf{s}) = \left\{ \mathbf{x} \in \mathbb{Z}^d \mid \exists \mathbf{z} \in \mathbb{Z}^e : \bigvee_{1 \leq i \leq v} A_i \mathbf{x} + B_i \mathbf{s} + D_i \mathbf{z} + \mathbf{c}_i \geq \mathbf{0} \right\}$$

is the set

$$H : \mathbb{Z}^n \rightarrow 2^{\mathbb{Z}^d} : \mathbf{s} \mapsto S(\mathbf{s}) = \left\{ \mathbf{x} \in \mathbb{Z}^d \mid \exists \mathbf{z} \in \mathbb{Z}^e : \bigwedge_{1 \leq i \leq v} A_i \mathbf{x} + B_i \mathbf{s} + D_i \mathbf{z} + \mathbf{c}_i + \mathbf{K}_i \geq \mathbf{0} \right\},$$

with \mathbf{K}_i the (component-wise) smallest non-negative integer vectors such that $S \subseteq H$.

The \mathbf{K}_i can be obtained by solving a number of LP problems, one for each element of each \mathbf{K}_i . If any LP problem is unbounded, then the corresponding constraint is dropped.

2.3 Coalescing

See Verdoolaege (2009), for now. More details will be added later.

2.4 Transitive Closure

2.4.1 Introduction

Definition 2.4.1 (Power of a Relation) Let $R \in \mathbb{Z}^n \rightarrow 2^{\mathbb{Z}^{d+d}}$ be a relation and $k \in \mathbb{Z}_{\geq 1}$ a positive number, then power k of relation R is defined as

$$R^k := \begin{cases} R & \text{if } k = 1 \\ R \circ R^{k-1} & \text{if } k \geq 2. \end{cases} \quad (2.1)$$

Definition 2.4.2 (Transitive Closure of a Relation) Let $R \in \mathbb{Z}^n \rightarrow 2^{\mathbb{Z}^{d+d}}$ be a relation, then the transitive closure R^+ of R is the union of all positive powers of R ,

$$R^+ := \bigcup_{k \geq 1} R^k.$$

Alternatively, the transitive closure may be defined inductively as

$$R^+ := R \cup (R \circ R^+). \quad (2.2)$$

Since the transitive closure of a polyhedral relation may no longer be a polyhedral relation (Kelly et al. 1996c), we can, in the general case, only compute an approximation of the transitive closure. Whereas Kelly et al. (1996c) compute underapproximations, we, like Beletskaya et al. (2009), compute overapproximations. That is, given a relation R , we will compute a relation T such that $R^+ \subseteq T$. Of course, we want this approximation to be as close as possible to the actual transitive closure R^+ and we want to detect the cases where the approximation is exact, i.e., where $T = R^+$.

For computing an approximation of the transitive closure of R , we follow the same general strategy as Beletskaya et al. (2009) and first compute an approximation of R^k for $k \geq 1$ and then project out the parameter k from the resulting relation.

Example 2.4.3 As a trivial example, consider the relation $R = \{x \rightarrow x + 1\}$. The k th power of this map for arbitrary k is

$$R^k = k \mapsto \{x \rightarrow x + k \mid k \geq 1\}.$$

The transitive closure is then

$$\begin{aligned} R^+ &= \{x \rightarrow y \mid \exists k \in \mathbb{Z}_{\geq 1} : y = x + k\} \\ &= \{x \rightarrow y \mid y \geq x + 1\}. \end{aligned}$$

2.4.2 Computing an Approximation of R^k

There are some special cases where the computation of R^k is very easy. One such case is that where R does not compose with itself, i.e., $R \circ R = \emptyset$ or $\text{dom } R \cap \text{ran } R = \emptyset$. In this case, R^k is only non-empty for $k = 1$ where it is equal to R itself.

In general, it is impossible to construct a closed form of R^k as a polyhedral relation. We will therefore need to make some approximations. As a first approximations, we will consider each of the basic relations in R as simply adding one or more offsets to a domain element to arrive at an image element and ignore the fact that some of these offsets may only be applied to some of the domain elements. That is, we will only consider the difference set ΔR of the relation. In particular, we will first construct a collection P of paths that move through a total of k offsets and then intersect domain and range of this collection with those of R . That is,

$$K = P \cap (\text{dom } R \rightarrow \text{ran } R), \quad (2.3)$$

with

$$P = \mathbf{s} \mapsto \{ \mathbf{x} \rightarrow \mathbf{y} \mid \exists k_i \in \mathbb{Z}_{\geq 0}, \delta_i \in k_i \Delta_i(\mathbf{s}) : \mathbf{y} = \mathbf{x} + \sum_i \delta_i \wedge \sum_i k_i = k > 0 \} \quad (2.4)$$

and with Δ_i the basic sets that compose the difference set ΔR . Note that the number of basic sets Δ_i need not be the same as the number of basic relations in R . Also note that since addition is commutative, it does not matter in which order we add the offsets and so we are allowed to group them as we did in (2.4).

If all the Δ_i s are singleton sets $\Delta_i = \{ \delta_i \}$ with $\delta_i \in \mathbb{Z}^d$, then (2.4) simplifies to

$$P = \{ \mathbf{x} \rightarrow \mathbf{y} \mid \exists k_i \in \mathbb{Z}_{\geq 0} : \mathbf{y} = \mathbf{x} + \sum_i k_i \delta_i \wedge \sum_i k_i = k > 0 \} \quad (2.5)$$

and then the approximation computed in (2.3) is essentially the same as that of Beletskaya et al. (2009). If some of the Δ_i s are not singleton sets or if some of δ_i s are parametric, then we need to resort to further approximations.

To ease both the exposition and the implementation, we will for the remainder of this section work with extended offsets $\Delta'_i = \Delta_i \times \{1\}$. That is, each offset is extended with an extra coordinate that is set equal to one. The paths constructed by summing such extended offsets have the length encoded as the difference of their final coordinates. The path P' can then be decomposed into paths P'_i , one for each Δ_i ,

$$P' = ((P'_m \cup \text{Id}) \circ \dots \circ (P'_2 \cup \text{Id}) \circ (P'_1 \cup \text{Id})) \cap \{ \mathbf{x}' \rightarrow \mathbf{y}' \mid y_{d+1} - x_{d+1} = k > 0 \}, \quad (2.6)$$

with

$$P'_i = \mathbf{s} \mapsto \{ \mathbf{x}' \rightarrow \mathbf{y}' \mid \exists k \in \mathbb{Z}_{\geq 1}, \delta \in k \Delta'_i(\mathbf{s}) : \mathbf{y}' = \mathbf{x}' + \delta \}.$$

Note that each P'_i contains paths of length at least one. We therefore need to take the union with the identity relation when composing the P'_i s to allow for paths that do not contain any offsets from one or more Δ'_i . The path that consists of only identity relations is removed by imposing the constraint $y_{d+1} - x_{d+1} > 0$. Taking the union with the identity relation means that that the relations we compose in (2.6) each consist of two

basic relations. If there are m disjuncts in the input relation, then a direct application of the composition operation may therefore result in a relation with 2^m disjuncts, which is prohibitively expensive. It is therefore crucial to apply coalescing (Section 2.3) after each composition.

Let us now consider how to compute an overapproximation of P'_i . Those that correspond to singleton Δ_i s are grouped together and handled as in (2.5). Note that this is just an optimization. The procedure described below would produce results that are at least as accurate. For simplicity, we first assume that no constraint in Δ'_i involves any existentially quantified variables. We will return to existentially quantified variables at the end of this section. Without existentially quantified variables, we can classify the constraints of Δ'_i as follows

1. non-parametric constraints

$$A_1 \mathbf{x} + \mathbf{c}_1 \geq \mathbf{0} \quad (2.7)$$

2. purely parametric constraints

$$B_2 \mathbf{s} + \mathbf{c}_2 \geq \mathbf{0} \quad (2.8)$$

3. negative mixed constraints

$$A_3 \mathbf{x} + B_3 \mathbf{s} + \mathbf{c}_3 \geq \mathbf{0} \quad (2.9)$$

such that for each row j and for all \mathbf{s} ,

$$\Delta'_i(\mathbf{s}) \cap \{ \delta' \mid B_{3,j} \mathbf{s} + c_{3,j} > 0 \} = \emptyset$$

4. positive mixed constraints

$$A_4 \mathbf{x} + B_4 \mathbf{s} + \mathbf{c}_4 \geq \mathbf{0}$$

such that for each row j , there is at least one \mathbf{s} such that

$$\Delta'_i(\mathbf{s}) \cap \{ \delta' \mid B_{4,j} \mathbf{s} + c_{4,j} > 0 \} \neq \emptyset$$

We will use the following approximation Q_i for P'_i :

$$Q_i = \mathbf{s} \mapsto \{ \mathbf{x}' \rightarrow \mathbf{y}' \mid \exists k \in \mathbb{Z}_{\geq 1}, \mathbf{f} \in \mathbb{Z}^d : \mathbf{y}' = \mathbf{x}' + (\mathbf{f}, k) \wedge A_1 \mathbf{f} + k \mathbf{c}_1 \geq \mathbf{0} \wedge B_2 \mathbf{s} + \mathbf{c}_2 \geq \mathbf{0} \wedge A_3 \mathbf{f} + B_3 \mathbf{s} + \mathbf{c}_3 \geq \mathbf{0} \}. \quad (2.10)$$

To prove that Q_i is indeed an overapproximation of P'_i , we need to show that for every $\mathbf{s} \in \mathbb{Z}^n$, for every $k \in \mathbb{Z}_{\geq 1}$ and for every $\mathbf{f} \in k \Delta_i(\mathbf{s})$ we have that (\mathbf{f}, k) satisfies the constraints in (2.10). If $\Delta_i(\mathbf{s})$ is non-empty, then \mathbf{s} must satisfy the constraints in (2.8). Each element $(\mathbf{f}, k) \in k \Delta'_i(\mathbf{s})$ is a sum of k elements $(\mathbf{f}_j, 1)$ in $\Delta'_i(\mathbf{s})$. Each of these elements satisfies the constraints in (2.7), i.e.,

$$\begin{bmatrix} A_1 & \mathbf{c}_1 \end{bmatrix} \begin{bmatrix} \mathbf{f}_j \\ 1 \end{bmatrix} \geq \mathbf{0}.$$

The sum of these elements therefore satisfies the same set of inequalities, i.e., $A_1\mathbf{f} + k\mathbf{c}_1 \geq \mathbf{0}$. Finally, the constraints in (2.9) are such that for any \mathbf{s} in the parameter domain of Δ , we have $-\mathbf{r}(\mathbf{s}) := B_3\mathbf{s} + \mathbf{c}_3 \leq \mathbf{0}$, i.e., $A_3\mathbf{f}_j \geq \mathbf{r}(\mathbf{s}) \geq \mathbf{0}$ and therefore also $A_3\mathbf{f} \geq \mathbf{r}(\mathbf{s})$. Note that if there are no mixed constraints and if the rational relaxation of $\Delta_i(\mathbf{s})$, i.e., $\{\mathbf{x} \in \mathbb{Q}^d \mid A_1\mathbf{x} + \mathbf{c}_1 \geq \mathbf{0}\}$, has integer vertices, then the approximation is exact, i.e., $Q_i = P'_i$. In this case, the vertices of $\Delta'_i(\mathbf{s})$ generate the rational cone $\{\mathbf{x}' \in \mathbb{Q}^{d+1} \mid [A_1 \quad \mathbf{c}_1]\mathbf{x}' \geq \mathbf{0}\}$ and therefore $\Delta'_i(\mathbf{s})$ is a Hilbert basis of this cone (Schrijver 1986, Theorem 16.4).

Existentially quantified variables can be handled by classifying them into variables that are uniquely determined by the parameters, variables that are independent of the parameters and others. The first set can be treated as parameters and the second as variables. Constraints involving the other existentially quantified variables are removed.

Example 2.4.4 Consider the relation

$$R = n \rightarrow \{x \rightarrow y \mid \exists \alpha_0, \alpha_1 : 7\alpha_0 = -2 + n \wedge 5\alpha_1 = -1 - x + y \wedge y \geq 6 + x\}.$$

The difference set of this relation is

$$\Delta = \Delta R = n \rightarrow \{x \mid \exists \alpha_0, \alpha_1 : 7\alpha_0 = -2 + n \wedge 5\alpha_1 = -1 + x \wedge x \geq 6\}.$$

The existentially quantified variables can be defined in terms of the parameters and variables as

$$\alpha_0 = \left\lfloor \frac{-2 + n}{7} \right\rfloor \quad \text{and} \quad \alpha_1 = \left\lfloor \frac{-1 + x}{5} \right\rfloor.$$

α_0 can therefore be treated as a parameter, while α_1 can be treated as a variable. This in turn means that $7\alpha_0 = -2 + n$ can be treated as a purely parametric constraint, while the other two constraints are non-parametric. The corresponding Q (2.10) is therefore

$$n \rightarrow \{(x, z) \rightarrow (y, w) \mid \exists \alpha_0, \alpha_1, k, f : k \geq 1 \wedge y = x + f \wedge w = z + k \wedge 7\alpha_0 = -2 + n \wedge 5\alpha_1 = -k + x \wedge x \geq 6k\}.$$

Projecting out the final coordinates encoding the length of the paths, results in the exact transitive closure

$$R^+ = n \rightarrow \{x \rightarrow y \mid \exists \alpha_0, \alpha_1 : 7\alpha_1 = -2 + n \wedge 6\alpha_0 \geq -x + y \wedge 5\alpha_0 \leq -1 - x + y\}.$$

2.4.3 Checking Exactness

The approximation T for the transitive closure R^+ can be obtained by projecting out the parameter k from the approximation K (2.3) of the power R^k . Since K is an overapproximation of R^k , T will also be an overapproximation of R^+ . To check whether the results are exact, we need to consider two cases depending on whether R is *cyclic*, where R is defined to be cyclic if R^+ maps any element to itself, i.e., $R^+ \cap \text{Id} \neq \emptyset$. If R is acyclic, then the inductive definition of (2.2) is equivalent to its completion, i.e.,

$$R^+ = R \cup (R \circ R^+)$$

is a defining property. Since T is known to be an overapproximation, we only need to check whether

$$T \subseteq R \cup (R \circ T).$$

This is essentially Theorem 5 of Kelly et al. (1996c). The only difference is that they only consider lexicographically forward relations, a special case of acyclic relations.

If, on the other hand, R is cyclic, then we have to resort to checking whether the approximation K of the power is exact. Note that T may be exact even if K is not exact, so the check is sound, but incomplete. To check exactness of the power, we simply need to check (2.1). Since again K is known to be an overapproximation, we only need to check whether

$$\begin{aligned} K'|_{y_{d+1}-x_{d+1}=1} &\subseteq R' \\ K'|_{y_{d+1}-x_{d+1} \geq 2} &\subseteq R' \circ K'|_{y_{d+1}-x_{d+1} \geq 1}, \end{aligned}$$

where $R' = \{ \mathbf{x}' \rightarrow \mathbf{y}' \mid \mathbf{x} \rightarrow \mathbf{y} \in R \wedge y_{d+1} - x_{d+1} = 1 \}$, i.e., R extended with path lengths equal to 1.

All that remains is to explain how to check the cyclicity of R . Note that the exactness on the power is always sound, even in the acyclic case, so we only need to be careful that we find all cyclic cases. Now, if R is cyclic, i.e., $R^+ \cap \text{Id} \neq \emptyset$, then, since T is an overapproximation of R^+ , also $T \cap \text{Id} \neq \emptyset$. This in turn means that $\Delta K'$ contains a point whose first d coordinates are zero and whose final coordinate is positive. In the implementation we currently perform this test on P' instead of K' . Note that if R^+ is acyclic and T is not, then the approximation is clearly not exact and the approximation of the power K will not be exact either.

2.4.4 Decomposing R into strongly connected components

If the input relation R is a union of several basic relations that can be partially ordered then the accuracy of the approximation may be improved by computing an approximation of each strongly connected components separately. For example, if $R = R_1 \cup R_2$ and $R_1 \circ R_2 = \emptyset$, then we know that any path that passes through R_2 cannot later pass through R_1 , i.e.,

$$R^+ = R_1^+ \cup R_2^+ \cup (R_2^+ \circ R_1^+). \quad (2.11)$$

We can therefore compute (approximations of) transitive closures of R_1 and R_2 separately. Note, however, that the condition $R_1 \circ R_2 = \emptyset$ is actually too strong. If $R_1 \circ R_2$ is a subset of $R_2 \circ R_1$ then we can reorder the segments in any path that moves through both R_1 and R_2 to first move through R_1 and then through R_2 .

This idea can be generalized to relations that are unions of more than two basic relations by constructing the strongly connected components in the graph with as vertices the basic relations and an edge between two basic relations R_i and R_j if R_i needs to follow R_j in some paths. That is, there is an edge from R_i to R_j iff

$$R_i \circ R_j \not\subseteq R_j \circ R_i. \quad (2.12)$$

The components can be obtained from the graph by applying Tarjan's algorithm (Tarjan 1972).

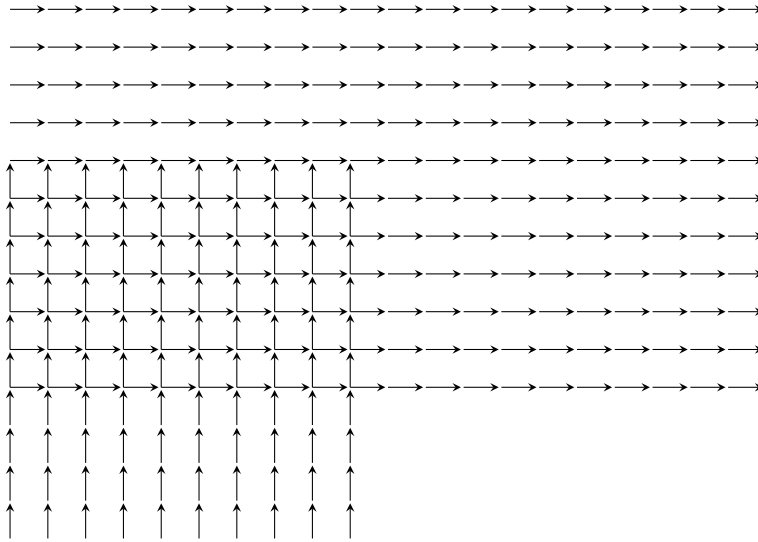


Figure 2.1: The relation from Example 2.4.5

In practice, we compute the (extended) powers K'_i of each component separately and then compose them as in (2.6). Note, however, that in this case the order in which we apply them is important and should correspond to a topological ordering of the strongly connected components. Simply applying Tarjan's algorithm will produce topologically sorted strongly connected components. The graph on which Tarjan's algorithm is applied is constructed on-the-fly. That is, whenever the algorithm checks if there is an edge between two vertices, we evaluate (2.12). The exactness check is performed on each component separately. If the approximation turns out to be inexact for any of the components, then the entire result is marked inexact and the exactness check is skipped on the components that still need to be handled.

It should be noted that (2.11) is only valid for exact transitive closures. If overapproximations are computed in the right hand side, then the result will still be an overapproximation of the left hand side, but this result may not be transitively closed. If we only separate components based on the condition $R_i \circ R_j = \emptyset$, then there is no problem, as this condition will still hold on the computed approximations of the transitive closures. If, however, we have exploited (2.12) during the decomposition and if the result turns out not to be exact, then we check whether the result is transitively closed. If not, we recompute the transitive closure, skipping the decomposition. Note that testing for transitive closedness on the result may be fairly expensive, so we may want to make this check configurable.

Example 2.4.5 Consider the relation in example `closure4` that comes with the `Omega`

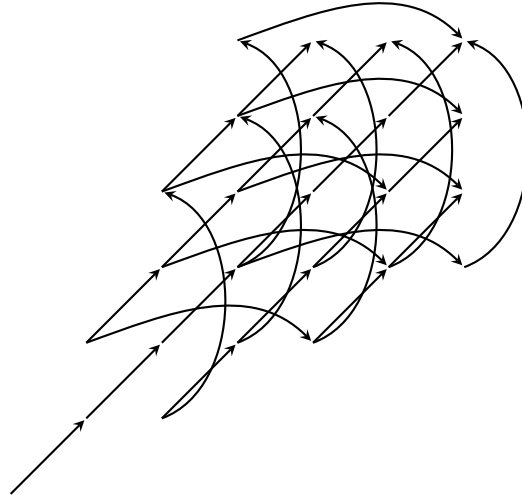


Figure 2.2: The relation from Example 2.4.6

calculator (Kelly et al. 1996a), $R = R_1 \cup R_2$, with

$$R_1 = \{(x, y) \rightarrow (x, y + 1) \mid 1 \leq x, y \leq 10\}$$

$$R_2 = \{(x, y) \rightarrow (x + 1, y) \mid 1 \leq x \leq 20 \wedge 5 \leq y \leq 15\}.$$

This relation is shown graphically in Figure 2.1. We have

$$R_1 \circ R_2 = \{(x, y) \rightarrow (x + 1, y + 1) \mid 1 \leq x \leq 9 \wedge 5 \leq y \leq 10\}$$

$$R_2 \circ R_1 = \{(x, y) \rightarrow (x + 1, y + 1) \mid 1 \leq x \leq 10 \wedge 4 \leq y \leq 10\}.$$

Clearly, $R_1 \circ R_2 \subseteq R_2 \circ R_1$ and so

$$(R_1 \cup R_2)^+ = (R_2^+ \circ R_1^+) \cup R_1^+ \cup R_2^+.$$

Example 2.4.6 Consider the relation on the right of Beletska et al. (2009, Figure 2), reproduced in Figure 2.2. The relation can be described as $R = R_1 \cup R_2 \cup R_3$, with

$$R_1 = n \mapsto \{(i, j) \rightarrow (i + 3, j) \mid i \leq 2j - 4 \wedge i \leq n - 3 \wedge j \leq 2i - 1 \wedge j \leq n\}$$

$$R_2 = n \mapsto \{(i, j) \rightarrow (i, j + 3) \mid i \leq 2j - 1 \wedge i \leq n \wedge j \leq 2i - 4 \wedge j \leq n - 3\}$$

$$R_3 = n \mapsto \{(i, j) \rightarrow (i + 1, j + 1) \mid i \leq 2j - 1 \wedge i \leq n - 1 \wedge j \leq 2i - 1 \wedge j \leq n - 1\}.$$

The figure shows this relation for $n = 7$. Both $R_3 \circ R_1 \subseteq R_1 \circ R_3$ and $R_3 \circ R_2 \subseteq R_2 \circ R_3$, which the reader can verify using the `iscc` calculator:

$$R1 := [n] \rightarrow \{ [i, j] \rightarrow [i+3, j] : i \leq 2j - 4 \text{ and } i \leq n - 3 \text{ and } j \leq 2i - 1 \text{ and } j \leq n \};$$

```

R2 := [n] -> { [i, j] -> [i, j+3] : i <= 2 j - 1 and i <= n and
                j <= 2 i - 4 and j <= n - 3 };
R3 := [n] -> { [i, j] -> [i+1, j+1] : i <= 2 j - 1 and i <= n - 1 and
                j <= 2 i - 1 and j <= n - 1 };
(R1 . R3) - (R3 . R1);
(R2 . R3) - (R3 . R2);

```

R_3 can therefore be moved forward in any path. For the other two basic relations, we have both $R_2 \circ R_1 \not\subseteq R_1 \circ R_2$ and $R_1 \circ R_2 \not\subseteq R_2 \circ R_1$ and so R_1 and R_2 form a strongly connected component. By computing the power of R_3 and $R_1 \cup R_2$ separately and composing the results, the power of R can be computed exactly using (2.5). As explained by Beletka et al. (2009), applying the same formula to R directly, without a decomposition, would result in an overapproximation of the power.

2.4.5 Partitioning the domains and ranges of R

The algorithm of Section 2.4.2 assumes that the input relation R can be treated as a union of translations. This is a reasonable assumption if R maps elements of a given abstract domain to the same domain. However, if R is a union of relations that map between different domains, then this assumption no longer holds. In particular, when an entire dependence graph is encoded in a single relation, as is done by, e.g., Barthou et al. (2000, Section 6.1), then it does not make sense to look at differences between iterations of different domains. Now, arguably, a modified Floyd-Warshall algorithm should be applied to the dependence graph, as advocated by Kelly et al. (1996c), with the transitive closure operation only being applied to relations from a given domain to itself. However, it is also possible to detect disjoint domains and ranges and to apply Floyd-Warshall internally.

Algorithm 1: The modified Floyd-Warshall algorithm of Kelly et al. (1996c)

Input: Relations R_{pq} , $0 \leq p, q < n$

Output: Updated relations R_{pq} such that each relation R_{pq} contains all indirect paths from p to q in the input graph

```

1 for  $r \in [0, n - 1]$  do
2    $R_{rr} := R_{rr}^+$ 
3   for  $p \in [0, n - 1]$  do
4     for  $q \in [0, n - 1]$  do
5       if  $p \neq r$  or  $q \neq r$  then
6          $R_{pq} := R_{pq} \cup (R_{rq} \circ R_{pr}) \cup (R_{rq} \circ R_{rr} \circ R_{pr})$ 

```

Let the input relation R be a union of m basic relations R_i . Let D_{2i} be the domains of R_i and D_{2i+1} the ranges of R_i . The first step is to group overlapping D_j until a partition is obtained. If the resulting partition consists of a single part, then we continue with the algorithm of Section 2.4.2. Otherwise, we apply Floyd-Warshall on the graph with as

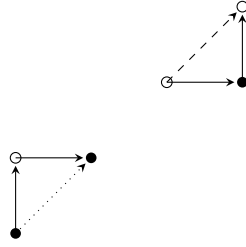


Figure 2.3: The relation (solid arrows) on the right of Figure 1 of Beletka et al. (2009) and its transitive closure

vertices the parts of the partition and as edges the R_i attached to the appropriate pairs of vertices. In particular, let there be n parts P_k in the partition. We construct n^2 relations

$$R_{pq} := \bigcup_{i \text{ s.t. } \text{dom } R_i \subseteq P_p \wedge \text{ran } R_i \subseteq P_q} R_i,$$

apply Algorithm 1 and return the union of all resulting R_{pq} as the transitive closure of R . Each iteration of the r -loop in Algorithm 1 updates all relations R_{pq} to include paths that go from p to r , possibly stay there for a while, and then go from r to q . Note that paths that “stay in r ” include all paths that pass through earlier vertices since R_{rr} itself has been updated accordingly in previous iterations of the outer loop. In principle, it would be sufficient to use the R_{pr} and R_{rq} computed in the previous iteration of the r -loop in Line 6. However, from an implementation perspective, it is easier to allow either or both of these to have been updated in the same iteration of the r -loop. This may result in duplicate paths, but these can usually be removed by coalescing (Section 2.3) the result of the union in Line 6, which should be done in any case. The transitive closure in Line 2 is performed using a recursive call. This recursive call includes the partitioning step, but the resulting partition will usually be a singleton. The result of the recursive call will either be exact or an overapproximation. The final result of Floyd-Warshall is therefore also exact or an overapproximation.

Example 2.4.7 Consider the relation on the right of Figure 1 of Beletka et al. (2009), reproduced in Figure 2.3. This relation can be described as

$$\{(x, y) \rightarrow (x_2, y_2) \mid (3y = 2x \wedge x_2 = x \wedge 3y_2 = 3 + 2x \wedge x \geq 0 \wedge x \leq 3) \vee (x_2 = 1 + x \wedge y_2 = y \wedge x \geq 0 \wedge 3y \geq 2 + 2x \wedge x \leq 2 \wedge 3y \leq 3 + 2x)\}.$$

Note that the domain of the upward relation overlaps with the range of the rightward relation and vice versa, but that the domain of neither relation overlaps with its own range or the domain of the other relation. The domains and ranges can therefore be partitioned into two parts, P_0 and P_1 , shown as the white and black dots in Figure 2.3,

respectively. Initially, we have

$$\begin{aligned} R_{00} &= \emptyset \\ R_{01} &= \{(x, y) \rightarrow (x + 1, y) \mid (x \geq 0 \wedge 3y \geq 2 + 2x \wedge x \leq 2 \wedge 3y \leq 3 + 2x)\} \\ R_{10} &= \{(x, y) \rightarrow (x_2, y_2) \mid (3y = 2x \wedge x_2 = x \wedge 3y_2 = 3 + 2x \wedge x \geq 0 \wedge x \leq 3)\} \\ R_{11} &= \emptyset. \end{aligned}$$

In the first iteration, R_{00} remains the same ($\emptyset^+ = \emptyset$). R_{01} and R_{10} are therefore also unaffected, but R_{11} is updated to include $R_{01} \circ R_{10}$, i.e., the dashed arrow in the figure. This new R_{11} is obviously transitively closed, so it is not changed in the second iteration and it does not have an effect on R_{01} and R_{10} . However, R_{00} is updated to include $R_{10} \circ R_{01}$, i.e., the dotted arrow in the figure. The transitive closure of the original relation is then equal to $R_{00} \cup R_{01} \cup R_{10} \cup R_{11}$.

2.4.6 Incremental Computation

In some cases it is possible and useful to compute the transitive closure of union of basic relations incrementally. In particular, if R is a union of m basic maps,

$$R = \bigcup_j R_j,$$

then we can pick some R_i and compute the transitive closure of R as

$$R^+ = R_i^+ \cup \left(\bigcup_{j \neq i} R_i^* \circ R_j \circ R_i^* \right)^+. \quad (2.13)$$

For this approach to be successful, it is crucial that each of the disjuncts in the argument of the second transitive closure in (2.13) be representable as a single basic relation, i.e., without a union. If this condition holds, then by using (2.13), the number of disjuncts in the argument of the transitive closure can be reduced by one. Now, $R_i^* = R_i^+ \cup \text{Id}$, but in some cases it is possible to relax the constraints of R_i^+ to include part of the identity relation, say on domain D . We will use the notation $C(R_i, D) = R_i^+ \cup \text{Id}_D$ to represent this relaxed version of R_i^+ . Kelly et al. (1996c) use the notation $R_i^?$. $C(R_i, D)$ can be computed by allowing k to attain the value 0 in (2.10) and by using

$$P \cap (D \rightarrow D)$$

instead of (2.3). Typically, D will be a strict superset of both $\text{dom } R_i$ and $\text{ran } R_i$. We therefore need to check that domain and range of the transitive closure are part of $C(R_i, D)$, i.e., the part that results from the paths of positive length ($k \geq 1$), are equal to the domain and range of R_i . If not, then the incremental approach cannot be applied for the given choice of R_i and D .

In order to be able to replace R^* by $C(R_i, D)$ in (2.13), D should be chosen to include both $\text{dom } R$ and $\text{ran } R$, i.e., such that $\text{Id}_D \circ R_j \circ \text{Id}_D = R_j$ for all $j \neq i$. Kelly et al. (1996c) say that they use $D = \text{dom } R_i \cup \text{ran } R_i$, but presumably they mean that they

use $D = \text{dom } R \cup \text{ran } R$. Now, this expression of D contains a union, so it not directly usable. Kelly et al. (1996c) do not explain how they avoid this union. Apparently, in their implementation, they are using the convex hull of $\text{dom } R \cup \text{ran } R$ or at least an approximation of this convex hull. We use the simple hull (Section 2.2) of $\text{dom } R \cup \text{ran } R$.

It is also possible to use a domain D that does *not* include $\text{dom } R \cup \text{ran } R$, but then we have to compose with $C(R_i, D)$ more selectively. In particular, if we have

$$\text{for each } j \neq i \text{ either } \text{dom } R_j \subseteq D \text{ or } \text{dom } R_j \cap \text{ran } R_i = \emptyset \quad (2.14)$$

and, similarly,

$$\text{for each } j \neq i \text{ either } \text{ran } R_j \subseteq D \text{ or } \text{ran } R_j \cap \text{dom } R_i = \emptyset \quad (2.15)$$

then we can refine (2.13) to

$$R_i^+ \cup \left(\left(\bigcup_{\substack{\text{dom } R_j \subseteq D \\ \text{ran } R_j \subseteq D}} C \circ R_j \circ C \right) \cup \left(\bigcup_{\substack{\text{dom } R_j \cap \text{ran } R_i = \emptyset \\ \text{ran } R_j \subseteq D}} C \circ R_j \right) \cup \left(\bigcup_{\substack{\text{dom } R_j \subseteq D \\ \text{ran } R_j \cap \text{dom } R_i = \emptyset}} R_j \circ C \right) \cup \left(\bigcup_{\substack{\text{dom } R_j \cap \text{ran } R_i = \emptyset \\ \text{ran } R_j \cap \text{dom } R_i = \emptyset}} R_j \right) \right)^+.$$

If only property (2.14) holds, we can use

$$R_i^+ \cup \left((R_i^+ \cup \text{Id}) \circ \left(\left(\bigcup_{\text{dom } R_j \subseteq D} R_j \circ C \right) \cup \left(\bigcup_{\text{dom } R_j \cap \text{ran } R_i = \emptyset} R_j \right) \right)^+ \right),$$

while if only property (2.15) holds, we can use

$$R_i^+ \cup \left(\left(\left(\bigcup_{\text{ran } R_j \subseteq D} C \circ R_j \right) \cup \left(\bigcup_{\text{ran } R_j \cap \text{dom } R_i = \emptyset} R_j \right) \right)^+ \circ (R_i^+ \cup \text{Id}) \right).$$

It should be noted that if we want the result of the incremental approach to be transitively closed, then we can only apply it if all of the transitive closure operations involved are exact. If, say, the second transitive closure in (2.13) contains extra elements, then the result does not necessarily contain the composition of these extra elements with powers of R_i .

2.4.7 An Omega-like implementation

While the main algorithm of Kelly et al. (1996c) is designed to compute and underapproximation of the transitive closure, the authors mention that they could also compute overapproximations. In this section, we describe our implementation of an algorithm that is based on their ideas. Note that the Omega library computes underapproximations (Kelly et al. 1996b, Section 6.4).

The main tool is Equation (2) of Kelly et al. (1996c). The input relation R is first overapproximated by a “d-form” relation

$$\{ \mathbf{i} \rightarrow \mathbf{j} \mid \exists \alpha : \mathbf{L} \leq \mathbf{j} - \mathbf{i} \leq \mathbf{U} \wedge (\forall p : j_p - i_p = M_p \alpha_p) \},$$

where p ranges over the dimensions and \mathbf{L} , \mathbf{U} and \mathbf{M} are constant integer vectors. The elements of \mathbf{U} may be ∞ , meaning that there is no upper bound corresponding to that element, and similarly for \mathbf{L} . Such an overapproximation can be obtained by computing strides, lower and upper bounds on the difference set ΔR . The transitive closure of such a “d-form” relation is

$$\{\mathbf{i} \rightarrow \mathbf{j} \mid \exists \alpha, k : k \geq 1 \wedge k\mathbf{L} \leq \mathbf{j} - \mathbf{i} \leq k\mathbf{U} \wedge (\forall p : j_p - i_p = M_p \alpha_p)\}. \quad (2.16)$$

The domain and range of this transitive closure are then intersected with those of the input relation. This is a special case of the algorithm in Section 2.4.2.

In their algorithm for computing lower bounds, the authors use the above algorithm as a substep on the disjuncts in the relation. At the end, they say

If an upper bound is required, it can be calculated in a manner similar to that of a single conjunct [sic] relation.

Presumably, the authors mean that a “d-form” approximation of the whole input relation should be used. However, the accuracy can be improved by also trying to apply the incremental technique from the same paper, which is explained in more detail in Section 2.4.6. In this case, $C(R_i, D)$ can be obtained by allowing the value zero for k in (2.16), i.e., by computing

$$\{\mathbf{i} \rightarrow \mathbf{j} \mid \exists \alpha, k : k \geq 0 \wedge k\mathbf{L} \leq \mathbf{j} - \mathbf{i} \leq k\mathbf{U} \wedge (\forall p : j_p - i_p = M_p \alpha_p)\}.$$

In our implementation we take as D the simple hull (Section 2.2) of $\text{dom } R \cup \text{ran } R$. To determine whether it is safe to use $C(R_i, D)$, we check the following conditions, as proposed by Kelly et al. (1996c): $C(R_i, D) - R_i^+$ is not a union and for each $j \neq i$ the condition

$$(C(R_i, D) - R_i^+) \circ R_j \circ (C(R_i, D) - R_i^+) = R_j$$

holds.

Bibliography

- Barthou, D., A. Cohen, and J.-F. Collard (2000). Maximal static expansion. *Int. J. Parallel Program.* 28(3), 213–243. [46]
- Beletska, A., D. Barthou, W. Bielecki, and A. Cohen (2009). Computing the transitive closure of a union of affine integer tuple relations. In *COCOA '09: Proceedings of the 3rd International Conference on Combinatorial Optimization and Applications*, Berlin, Heidelberg, pp. 98–109. Springer-Verlag. [39, 40, 45, 46, 47]
- Kelly, W., V. Maslov, W. Pugh, E. Rosser, T. Shpeisman, and D. Wonnacott (1996a, November). The Omega calculator and library. Technical report, University of Maryland. [45]
- Kelly, W., V. Maslov, W. Pugh, E. Rosser, T. Shpeisman, and D. Wonnacott (1996b, November). The Omega library. Technical report, University of Maryland. [49]
- Kelly, W., W. Pugh, E. Rosser, and T. Shpeisman (1996c). Transitive closure of infinite graphs and its applications. In C.-H. Huang, P. Sadayappan, U. Banerjee, D. Gelernter, A. Nicolau, and D. A. Padua (Eds.), *Languages and Compilers for Parallel Computing, 8th International Workshop, LCPC'95, Columbus, Ohio, USA, August 10-12, 1995, Proceedings*, Volume 1033 of *Lecture Notes in Computer Science*, pp. 126–140. Springer. [39, 43, 46, 48, 49, 50]
- Schrijver, A. (1986). *Theory of Linear and Integer Programming*. John Wiley & Sons. [42]
- Tarjan, R. (1972). Depth-first search and linear graph algorithms. *SIAM Journal on Computing* 1(2), 146–160. [43]
- Verdoolaege, S. (2009, April). An integer set library for program analysis. Advances in the Theory of Integer Linear Optimization and its Extensions, AMS 2009 Spring Western Section Meeting, San Francisco, California, 25-26 April 2009. [39]