

More Local and Semi-Local SETL Optimisations.

Optimisations may be divided into three classes:

A. *Local* optimisations, which can be performed when certain strictly local code features are detected.

B. *Semi-local* optimisations, which involve only local code transformation but which can only be performed when certain local 'code' features are detected and in addition one has verified that a program to be optimised possesses certain necessary global properties.

C. *Global* optimisations, which involve global program properties and make use of non-local code transformations.

In this newsletter, which continues D. Shields' Newsletter 45 on the same subject, we will enumerate various optimisations of type A and B. Note that optimisations of type A can be built directly into the new SETLB to BALM translator; optimisations of type B must wait until some degree of global program analysis is available in the SETL system.

To get full benefit even from the strictly local optimisations, they may have to be carried through all the levels of translation which the SETL system is now using. That is, the SETLB to BALM translator may notice a feature, and translate it using a specially introduced BALM operation, which the BALM to LITTLE translator is able to detect and convert into an appropriately efficient LITTLE sequence. As an example of this observation, note that the n appearing in optimisation (iii) below is known to be an integer, and even a 'short' integer; to exploit this fact properly, it must be reflected even at the LITTLE level. A similar remark applies to optimisation (v) below.

I. Local optimisations: (Some of which are quoted from Newsletter 45).

- i. Compile $\exists\{e(x), x \in a \mid C(x)\}$ as
if $\exists x \in a \mid C(x)$ then $e(x)$ else Ω
- ii. Compile $y \in \{e(x), x \in a \mid C(x)\}$ as
temp = f; ($\forall x \in a \mid C(x)$) if y eq e(x) then temp = t; quit;;
- iii. Compile $\#\{x \in a \mid C(x)\}$ as
n = 0; ($\forall x \in a \mid C(x)$) n = n + 1;;

(note that the SRTL level *next* function can be used explicitly)

iv. Calculation of all constant vectors and sets can be moved to an initialisation section; constants being assigned to global names. If possible, this should be done also for constants involving labels; if this is not possible, a 'once only' switch should be prefixed to the calculation of a constant which involves labels.

v. Special equality and inequality tests such as x eq nl, x ne nl, type x eq int etc. should be done directly by efficient BALM, or even by specially added LITTLE macros.

vi. The copy operation normally applied to the first argument of $s + t$ (if s is a set) can be omitted if s is an expression. Since '+' is commutative, this remark applies also if t is an expression. It also applies to $s - t$ if s is an expression.

vii. The union operation

$$s + \{e(x), x \in a \mid C(x)\}$$

can be compiled as

$$(\forall x \in a \mid C(x)) x \text{ in } s;;$$

viii. The intersection operation

$$s \cap \{e(x), x \in a \mid C(x)\}$$

can be compiled as

$$t = \text{nl}; (\forall x \in a \mid C(x)) \text{if } (e(x) \text{ is elt) \text{ es then elt in } t;;$$

ix. $\{x \in a \mid C(x)\}$ ne nl becomes $\exists x \in a \mid C(x)$

x. set - $\{e(x), x \in a \mid C(x)\}$ becomes

$$(\forall x \in a \mid C(x)) e(x) \text{ out } \text{set};;$$


```

body;
m1 = m1 + 1;
go to loop;

```

The same optimisation can be applied to descending iterations.

xiii. In calculating

$$[+: x \in s \mid C(x)] \langle e(x) \rangle,$$

pre-estimate the number of components of the resulting vector as being equal to # s; pre-allocate a vector of this size, and build up the required result by component insertion, finally throwing away excess space. More generally, when an expression of this form $[+: x \in s \mid C(x)] e(x)$ occurs in the context $[+: x \in s \mid C(x)] \langle e(x) \rangle + [+: x \in s_1 \mid C_1(x)] \langle e_1(x) \rangle$ calculation of the whole ought to be arranged to avoid the creating of unnecessary vector fragments and unnecessary copying.

xiv. Enhance the SETLB syntax so as to allow iterators over tuples, i.e. iterators $(\forall x(k) \in t) \dots$. Iterations of this sort can employ the optimisation described as (xii) above. Moreover, within an iterative loop headed by $(\forall x(k) \in t) \dots$, the indexed access $x(k)$ can be direct and efficient. The same optimisation should of course be applied to forms such as $\exists x(k) \in t, \{ e(x), x(k) \in t \mid \dots \}$, etc.

xv. Iterations of the form $(\forall x \in f\{s\})$ can be handled in a special way. Specifically, the first element address of a tuple whose first element is s can be located, and a routine called to reconstruct the tail of this tuple. Then the present *nextelt* routine can be used to advance the element address, and iteration can proceed until all tuples starting with s have been exhausted. Note that the set $f\{s\}$ never has to be constructed explicitly.

A similar approach can be applied to iterators $\forall x \in f\{s_1, s_2\}$, and to iterator-based constructions such as $\{e(x), x \in f\{s\} | C(x)\}$, etc.

II. Semi-Local Optimisations.

Some of the optimisations noted in this section depend on the availability of live-dead and object-type information. Of course, these are harder optimisations than those described in the preceding section.

i. Quantities known to be integers might be handled in a special way as follows:

Make available two compile time declarations NOERROR (permitting the suppression of certain costly types of run-time error checking, and also permitting the use of type information derived by A. Tenenbaum's 'backward' method) and SHORT (declaring that no integers in excess of an implementation-defined limit will occur.) In the presence of these declarations, integers can be maintained as LITTLE integers, and be converted to SETL form only when they are put in sets, become vector components, or assigned to a variable. For LITTLE integers, fast arithmetic comparison operations are available.

ii. Sets which never become set elements or vector components, and which are never assigned to more than one variable, will be dead when the sole variable x of which they are the value becomes dead. Value semantics can be guaranteed for such sets even without maintaining dynamic reference counts. When such sets s are used in constructions such as $s + s_1$, copying of s can be omitted. Similar remarks applies to tuples t appearing in the context $t + t_1$.

iii. In constructions like

$$s = \{e(x), x \in a \mid C(x)\};$$

...

$$t = s + \{e_1(x), x \in a_1 \mid C_1(x)\};$$

it may be possible to detect that s is used only to form a union, and hence to suppress the explicit construction of s . This same remark applies to sets used only to form intersections.