

Color Preferences



Sometimes we wish to assign a particular register (color) to a selected Live Range (e.g., a parameter or return value) *if possible*.

We can mark a node in the Interference Graph with a *Color Preference*.

When we unstack nodes and assign colors, we will avoid choosing color *c* if an uncolored neighbor has indicated a preference for it. If only color *c* is left, we take it (and ignore the preference).

Example

Assume in our previous example that *lim1* has requested register R1 and *lim2* has requested register R2 (because these are the registers the parameters are passed in).



Now when *i*, *j* and *t1* are unstacked, they respect *lim1*'s and *lim2*'s preferences:

i:R3, *j*:R1, *t2*:R2, *t1*:R2, *lim1*:R1, *lim2*:spill

Using Coloring to Optimize Register Moves

A nice "fringe benefit" of allocating registers via coloring is that we can often *optimize away* register to register moves by giving the source and target the *same color*.

Consider

Live in: *a*, *b*


t1 = *a* + *b*

x = *t1*

y = *x* + 1

q = *t1*

Live out: *y*, *q*



We'd like *x*, *t1* and *q* to get the same color. How do we "force" this?

We can “merge” x , $t1$ and q together:

Live in: a, b


$t1 = a + b$

$x = t1$

$y = x + 1$

$q = t1$

Live out: y, q




Now a 2-coloring that optimizes away both register to register moves is trivial.

Reckless Coalescing

Originally, Chaitin suggested merging *all* move-related nodes that don't interfere.

This is *reckless*—the merged node may not be colorable!

(Is it worth a spill to save a move??)



This Graph is 2-colorable before the reckless merge, but *not* after.


Conservative Coalescing

In response to Chaitin's reckless coalescing approach, Briggs suggested a *more conservative* approach.

See “Improvement to Graph Coloring Register Allocation,” P. Briggs et. al., ACM Toplas, May 1994.

Briggs suggested that two move-related nodes should be merged *only if* the combined source and target node has fewer than R neighbors.

This *guarantees* that the combined node will be colorable, but may miss some optimization opportunities.



After a merge of nodes a and d , there will be four neighbors, but a 2-coloring is still possible.

Iterated Coalescing

This is an intermediate approach, that seeks to be safer than reckless coalescing and more effective than conservative coalescing. It was proposed by George and Appel.

1. Build:

Create an Interference Graph, as usual. Mark source-target pairs with a special move-related arc (denoted as a dashed line).

2. Simplify:

Remove and stack non-move-related nodes with $< R$ neighbors.

3. Coalesce:

Combine move-related pairs that will have $< R$ neighbors after coalescing.

Repeat steps 2 and 3 until only nodes with R or more neighbors or move-related nodes remain or the graph is empty.

4. Freeze:

If the Interference Graph is non-empty:

Then If there exists a move-related node with $< R$ neighbors
Then: "Freeze in" the move and make the node non-move-related.

Return to Steps 2 and 3.


Else: Use Chaitin's Cost/Neighbors criterion to remove and stack a node.
Return to Steps 2 and 3.

5. Unstack:

Color nodes as they are unstacked as per Chaitin and Briggs.


Example

```
Live in: k,j
g = mem[j+12]
h = k-1
f = g*h
e = mem[j+8]
m = mem[j+16]
b = mem[f]
c = e+8
d = c
k = m+4
j = b
goto d
Live out: d,k,j
```




Assume we want a 4-coloring.

Note that neither j & b nor d & c can be conservatively colored.




We simplify by removing nodes with fewer than 4 neighbors.
 We remove and stack: **g, h, k, f, e, m**


The remaining Interference Graph is



We can now conservatively coalesce the move-related pairs to obtain



These remaining nodes can now be removed and stacked.



We can now unstack and color:
d&c:R1, j&b:R2, m:R3, e:R4, f:R1, k:R3, h:R1, g:R4

No spills were required and both moves were optimized away.

Reading Assignment

- Read David Wall's paper, "Global Register Allocation at Link Time."

Priority-Based Register Allocation

Alternatives to Chaitin-style register allocation are presented in:

- Hennessy and Chow, "The priority-based coloring approach to register allocation," ACM TOPLAS, October 1990.
- Larus and Hilfinger, "Register allocation in the SPUR Lisp compiler," SIGPLAN symposium on Compiler Construction, 1986.

These papers suggest two innovations:

1. Use of a *Priority Value* to choose nodes to color in an Interference Graph.

A Priority measures
 $(\text{Spill cost})/(\text{Size of Live Range})$

The idea is that small live ranges with a high spill cost are ideal candidates for register allocation.

As the size of a live range grows, it becomes less attractive for register allocation (since it "ties up" a register for a larger portion of a program).

2. Live Range Splitting

Rather than spill an entire live range that can't be colored, the live range is split into two or more smaller live ranges that may be colorable.

Large vs. Small Live Ranges

- A large live range has less spill code. Values are directly read from and written to a register.
But, a large live range is harder to allocate, since it may conflict with many other register candidates.
- A small live range is easier to allocate since it competes with fewer register candidates.
But, more spill code is needed to load and save register values across live ranges.
- In the limit a live range can shrink to a single definition or use of a register.
But, then we really don't have an effective register allocation at all!

Terminology

In an Interference Graph:

- A node with fewer neighbors than colors is termed *unconstrained*. It is trivial to color.
- A node that is not unconstrained is termed *constrained*. It may need to be split or spilled.