1. Cache Optimization:

Idea: local memory access

- Loading from memory fetches not just one byte but fills a complete cache line.
- Access to neighbored cells become cheaper.
- If all data of an inner loop fits into the cache, the iteration becomes maximally memory-efficient ...

Possible Solutions:

$\rightarrow$ Reorganize the data accesses !

$\rightarrow$ Reorganize the data !

Such optimizations can be made fully automatic only for arrays :-(

Example:

$$
\begin{aligned}
&\textsf{for}\ \ (j = 1; j < n; j{+}{+}) \\
&\qquad \textsf{for}\ \ (i = 1; i < m; i{+}{+}) \\
&\qquad\qquad a[i][j] = a[i-1][j-1] + a[i][j];
\end{aligned}
$$

$\Longrightarrow$ At first, always iterate over the rows!

$\Longrightarrow$ Exchange the ordering of the iterations:

$$
\begin{aligned}
&\textsf{for}\ \ (i = 1; i < m; i{+}{+}) \\
&\qquad \textsf{for}\ \ (j = 1; j < n; j{+}{+}) \\
&\qquad\qquad a[i][j] = a[i-1][j-1] + a[i][j];
\end{aligned}
$$

When is this permitted???

Iteration Scheme: before:

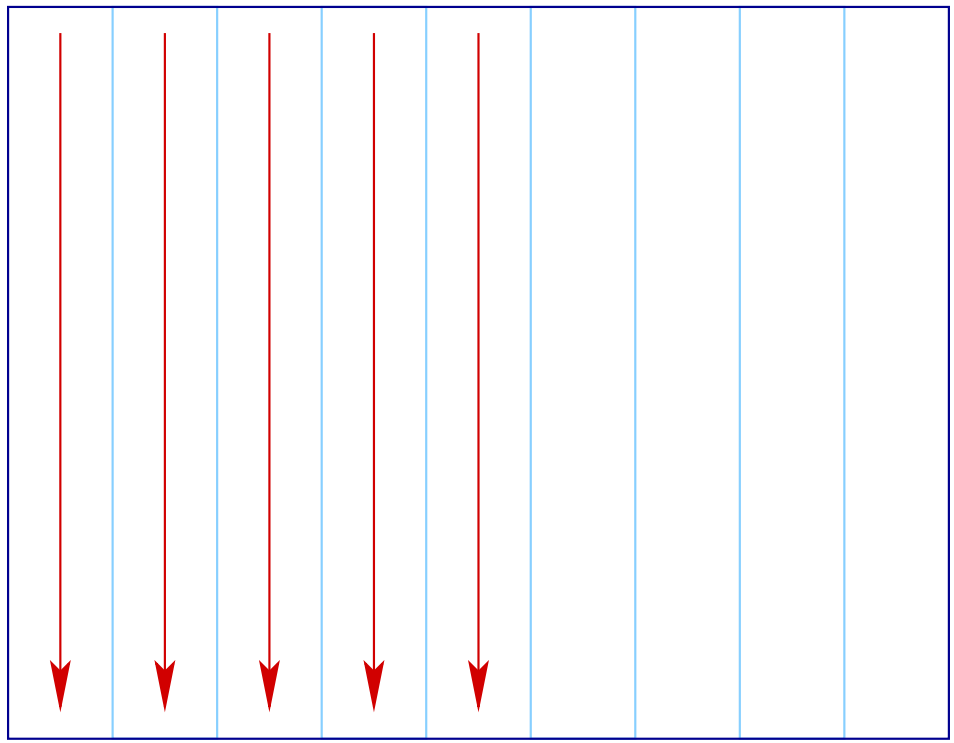

Iteration Scheme: after:

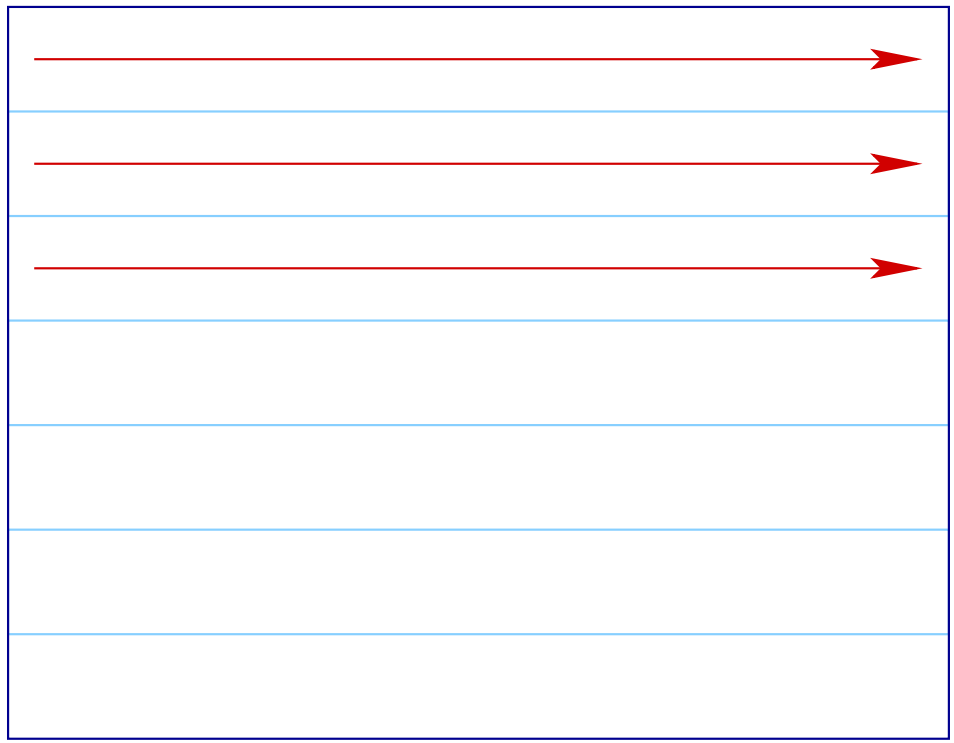

Iteration Scheme: allowed dependencies:

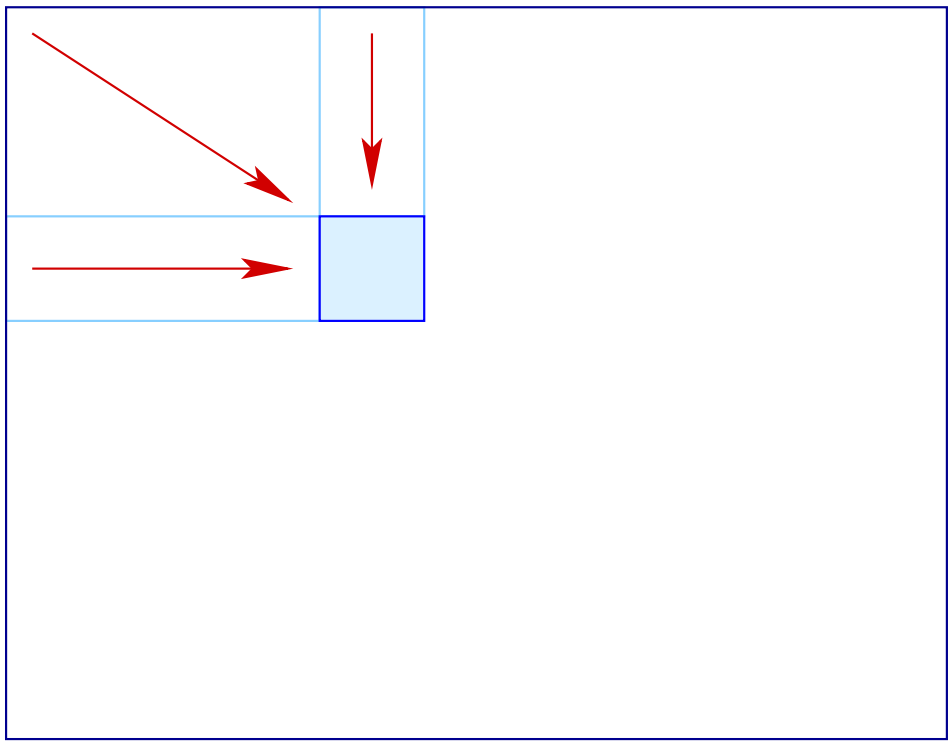

In our case, we must check that the following equation systems have no solution:

Write		Read
(i_1, j_1)	$=$	$(i_2 - 1, j_2 - 1)$
i_1	$\leq$	i_2
j_2	$\leq$	j_1
(i_1, j_1)	$=$	$(i_2 - 1, j_2 - 1)$
i_2	$\leq$	i_1
j_1	$\leq$	j_2

The first implies: $j_2 \leq j_2 - 1$ Hurra!

The second implies: $i_2 \leq i_2 - 1$ Hurra!

Example: Matrix-Matrix Multiplication

$$
\begin{aligned}
&\textsf{for}\ \ (i = 0; i < N; i{+}{+}) \\
&\qquad \textsf{for}\ \ (j = 0; j < M; j{+}{+}) \\
&\qquad\qquad \textsf{for}\ \ (k = 0; k < K; k{+}{+}) \\
&\qquad\qquad\qquad c[i][j] = c[i][j] + a[i][k] \cdot b[k][j];
\end{aligned}
$$

Over $b[][]$ the iteration is columnwise :-(

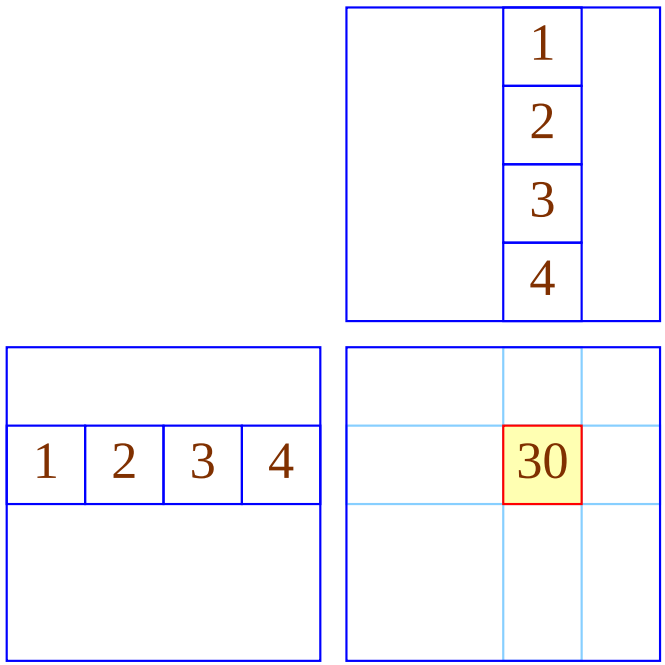
1
2
3
4
1
2
3
4
30

Exchange the two inner loops:

$$
\begin{aligned}
&\textsf{for}\ \ (i = 0; i < N; i{+}{+}) \\
&\qquad \textsf{for}\ \ (k = 0; k < K; k{+}{+}) \\
&\qquad\qquad \textsf{for}\ \ (j = 0; j < M; j{+}{+}) \\
&\qquad\qquad\qquad c[i][j] = c[i][j] + a[i][k] \cdot b[k][j];
\end{aligned}
$$

Is this permitted ???

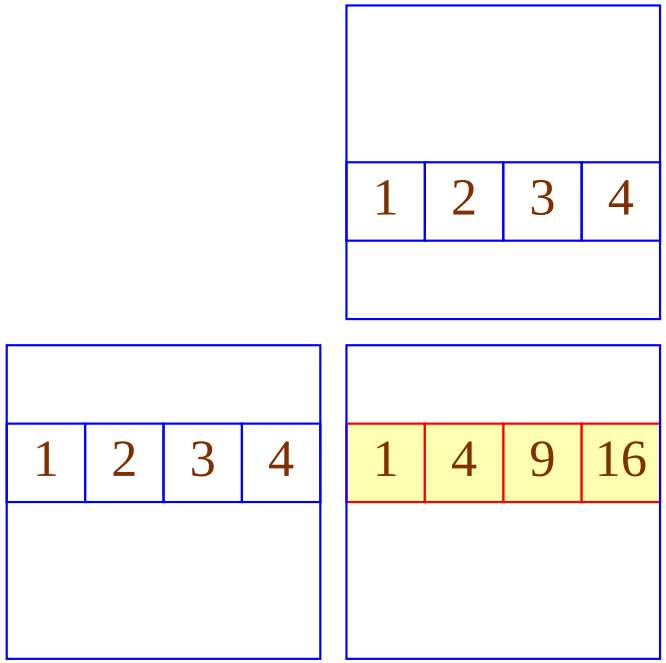
1 2 3 4
1 2 3 4
1 4 9 16

Discussion:

- Correctness follows as before :-)
- A similar idea can also be used for the implementation of multiplication for row compressed matrices :-))
- Sometimes, the program must be massaged such that the transformation becomes applicable :-(
- Matrix-matrix multiplication perhaps requires initialization of the result matrix first ...

$$
\begin{array}{l}
\textsf{for}\ \ (i = 0; i < N; i{+}{+}) \\
\quad \textsf{for}\ \ (j = 0; j < M; j{+}{+})\ \ \{ \\
\qquad c[i][j] = 0; \\
\qquad \textsf{for}\ \ (k = 0; k < K; k{+}{+}) \\
\qquad\quad c[i][j] = c[i][j] + a[i][k] \cdot b[k][j]; \\
\quad \}
\end{array}
$$

- Now, the two iterations can no longer be exchanged :-(
- The iteration over j, however, can be duplicated ...

$$
\begin{array}{l}
\textsf{for}\ \ (i = 0; i < N; i{+}{+})\ \ \{ \\
\qquad \textsf{for}\ \ (j = 0; j < M; j{+}{+})\ \ c[i][j] = 0; \\
\qquad \textsf{for}\ \ (j = 0; j < M; j{+}{+}) \\
\qquad\qquad \textsf{for}\ \ (k = 0; k < K; k{+}{+}) \\
\qquad\qquad\qquad c[i][j] = c[i][j] + a[i][k] \cdot b[k][j]; \\
\}
\end{array}
$$

Correctness:

$\Longrightarrow$ The read entries (here: no) may not be modified in the remaining body of the loop !!!

$\Longrightarrow$ The ordering of the write accesses to a memory cell may not be changed :-)

We obtain:

$$
\begin{array}{l}
\textsf{for}\ \ (i = 0; i < N; i{+}{+})\ \ \{ \\
\qquad \textsf{for}\ \ (j = 0; j < M; j{+}{+})\ \ c[i][j] = 0; \\
\qquad \textsf{for}\ \ (k = 0; k < K; k{+}{+}) \\
\qquad\qquad \textsf{for}\ \ (j = 0; j < M; j{+}{+}) \\
\qquad\qquad\qquad c[i][j] = c[i][j] + a[i][k] \cdot b[k][j]; \\
\}
\end{array}
$$

Discussion:

- Instead of fusing several loops, we now have distributed the loops :-)
- Accordingly, conditionals may be moved out of the loop $\Longrightarrow$ if-distribution ...

Warning:

Instead of using this transformation, the inner loop could also be optimized as follows:

$$
\begin{array}{l}
\textsf{for}\ \ (i = 0; i < N; i{+}{+}) \\
\qquad \textsf{for}\ \ (j = 0; j < M; j{+}{+})\ \ \{ \\
\qquad\qquad t = 0; \\
\qquad\qquad \textsf{for}\ \ (k = 0; k < K; k{+}{+}) \\
\qquad\qquad\qquad\quad t = t + a[i][k] \cdot b[k][j]; \\
\qquad\qquad c[i][j] = t; \\
\qquad \}
\end{array}
$$

Idea:

If we find heavily used array elements $a[e_1] \ldots [e_r]$ whose index expressions stay constant within the inner loop, we could instead also provide auxiliary registers :-)

Warning:

The latter optimization prohibits the former and vice versa ...

Discussion:

- so far, the optimizations are concerned with iterations over arrays.
- Cache-aware organization of other data-structures is possible, but in general not fully automatic ...

Example: Stacks

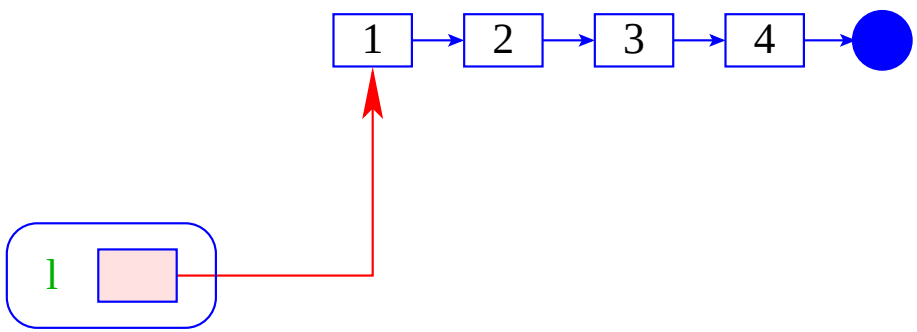

Advantage:

+ The implementation is simple :-)
+ The operations push / pop require constant time :-)
+ The data-structure may grow arbitrarily :-)

Disadvantage:

– The individual list objects may be arbitrarily dispersed over the memory :-(

Alternative:

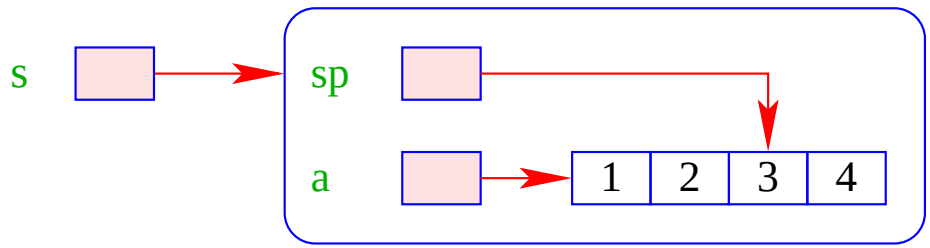

Advantage:

+ The implementation is also simple :-)
+ The operations push / pop still require constant time :-)
+ The data are consequtively allocated; stack oscillations are typically small

$\Longrightarrow$ better Cache behavior !!!

Disadvantage:

– The data-structure is bounded :-(

Improvement:

- If the array is full, replace it with another of double size !!!
- If the array drops empty to a quarter, halve the array again !!!

$\Longrightarrow$ The extra amortized costs are constant :-)

$\Longrightarrow$ The implementation is no longer so trivial :-}

Discussion:

$\rightarrow$ The same idea also works for queues :-)

$\rightarrow$ Other data-structures are attempted to organize blockwise.

Problem: how can accesses be organized such that they refer mostly to the same block ???

$\Longrightarrow$ Algorithms for external data