

Overview

- Floating-point arithmetic suffers from rounding error
- Tradeoff between faster single precision (32 bit) and more accurate double precision (64 bit)
- We have built a prototype system to auto-tune the precision level of particular components of a program
- Initial results with the Algebraic MultiGrid kernel demonstrate a nearly 2X speedup

Background

- **Mixed precision:** execute a few crucial instructions in double precision and the rest in single precision

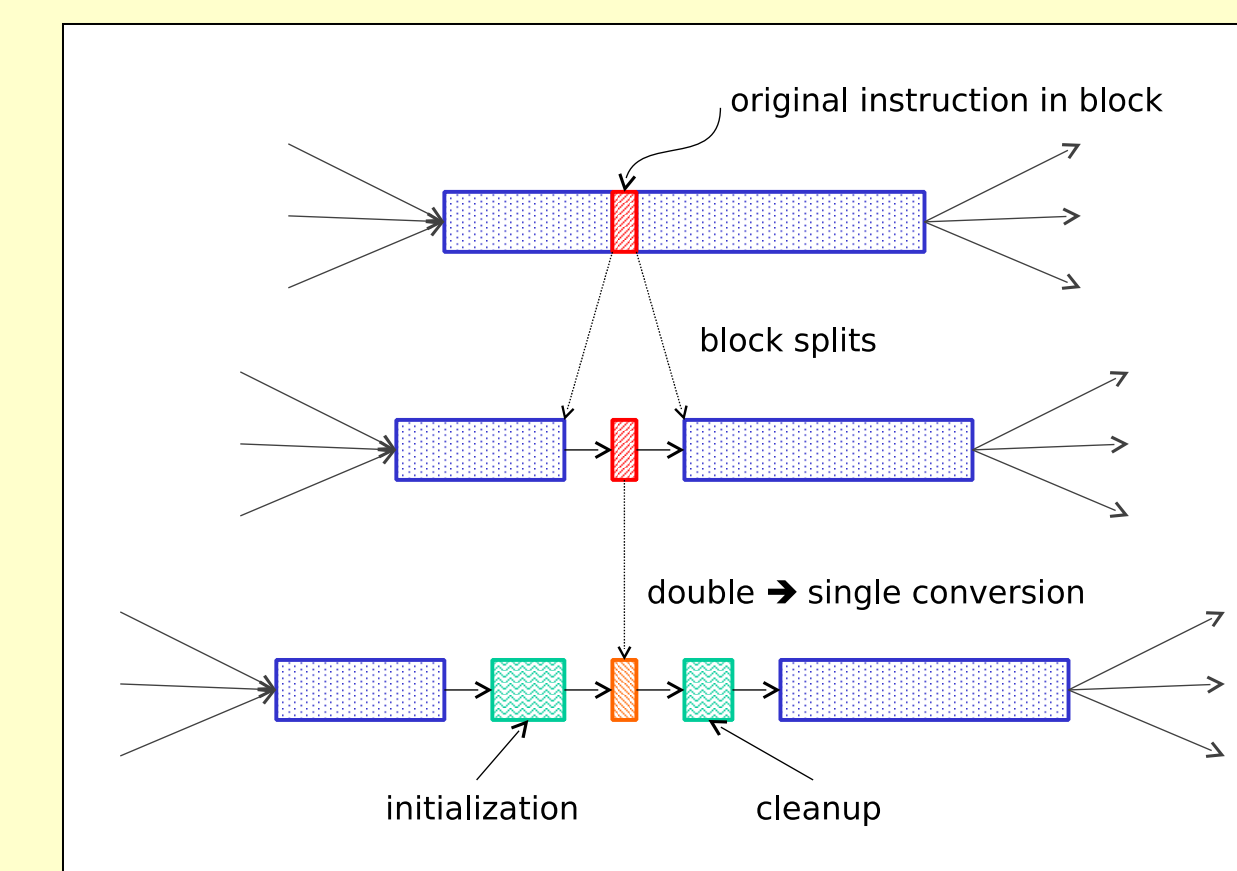
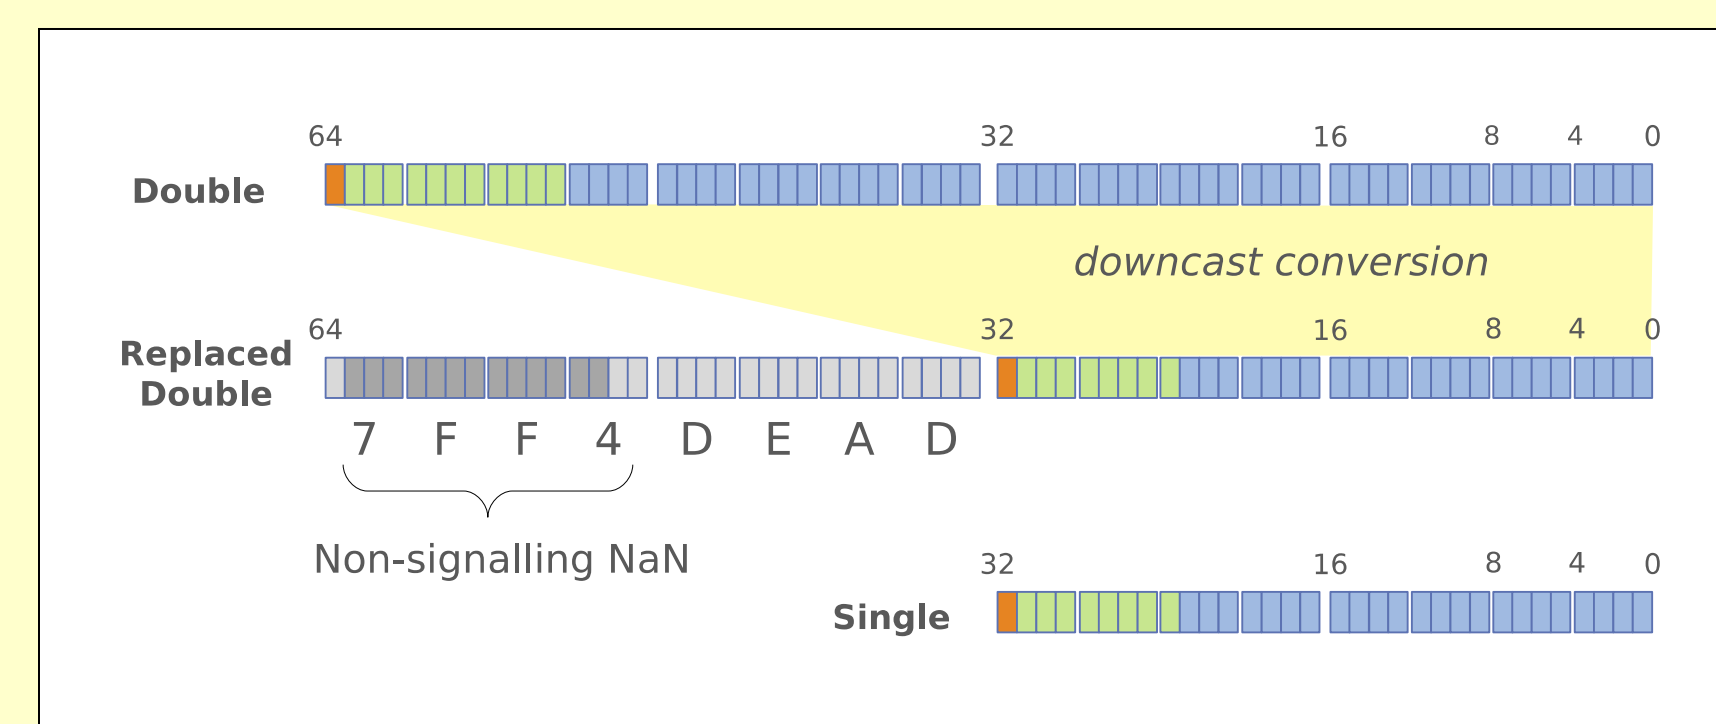
```

1: LU ← PA
2: solve Ly=Pb
3: solve Ux0 = y
4: for k = 1,2,... do
5:     rk ← b - Axk-1 (*)
6:     solve Ly = Prk
7:     solve Uzk = y
8:     xk ← xk-1 + zk (*)
9:     convergence check
10: end for
    
```

Double-precision operations in **red text** with astericks

Methods

- Replace double-precision instructions and values with single-precision versions
- Basic block patching (binary modification)



- Mixed-precision configuration representation

