


PERFORMANCE (REALLY) MATTERS

Emery Berger


UMASS AMHERST

(joint work with Charlie Curtsinger, Grinnell College)


emeryberger.com, [@emeryberger](https://twitter.com/emeryberger)



2020




2020




1980s

2020s



1980s



2020



2.65 GHz x 6

64-bit processor

4Gb RAM

1980s



4.77 MHz

16-bit processor

640Kb RAM



2.65 GHz x 6
64-bit processor
4Gb RAM



4.77 MHz
16-bit processor
640Kb RAM

2.65 GHz x 6

64-bit processor

4Gb RAM




≈ 3000x

4.77 MHz

16-bit processor

640Kb RAM

-1980s



40 Year Performance Ride!



1970

1975

1980

1985

1990

1995

2000

2005

2010

2015

Year

1970

1975

1980

1985

1990

1995

2000

2005

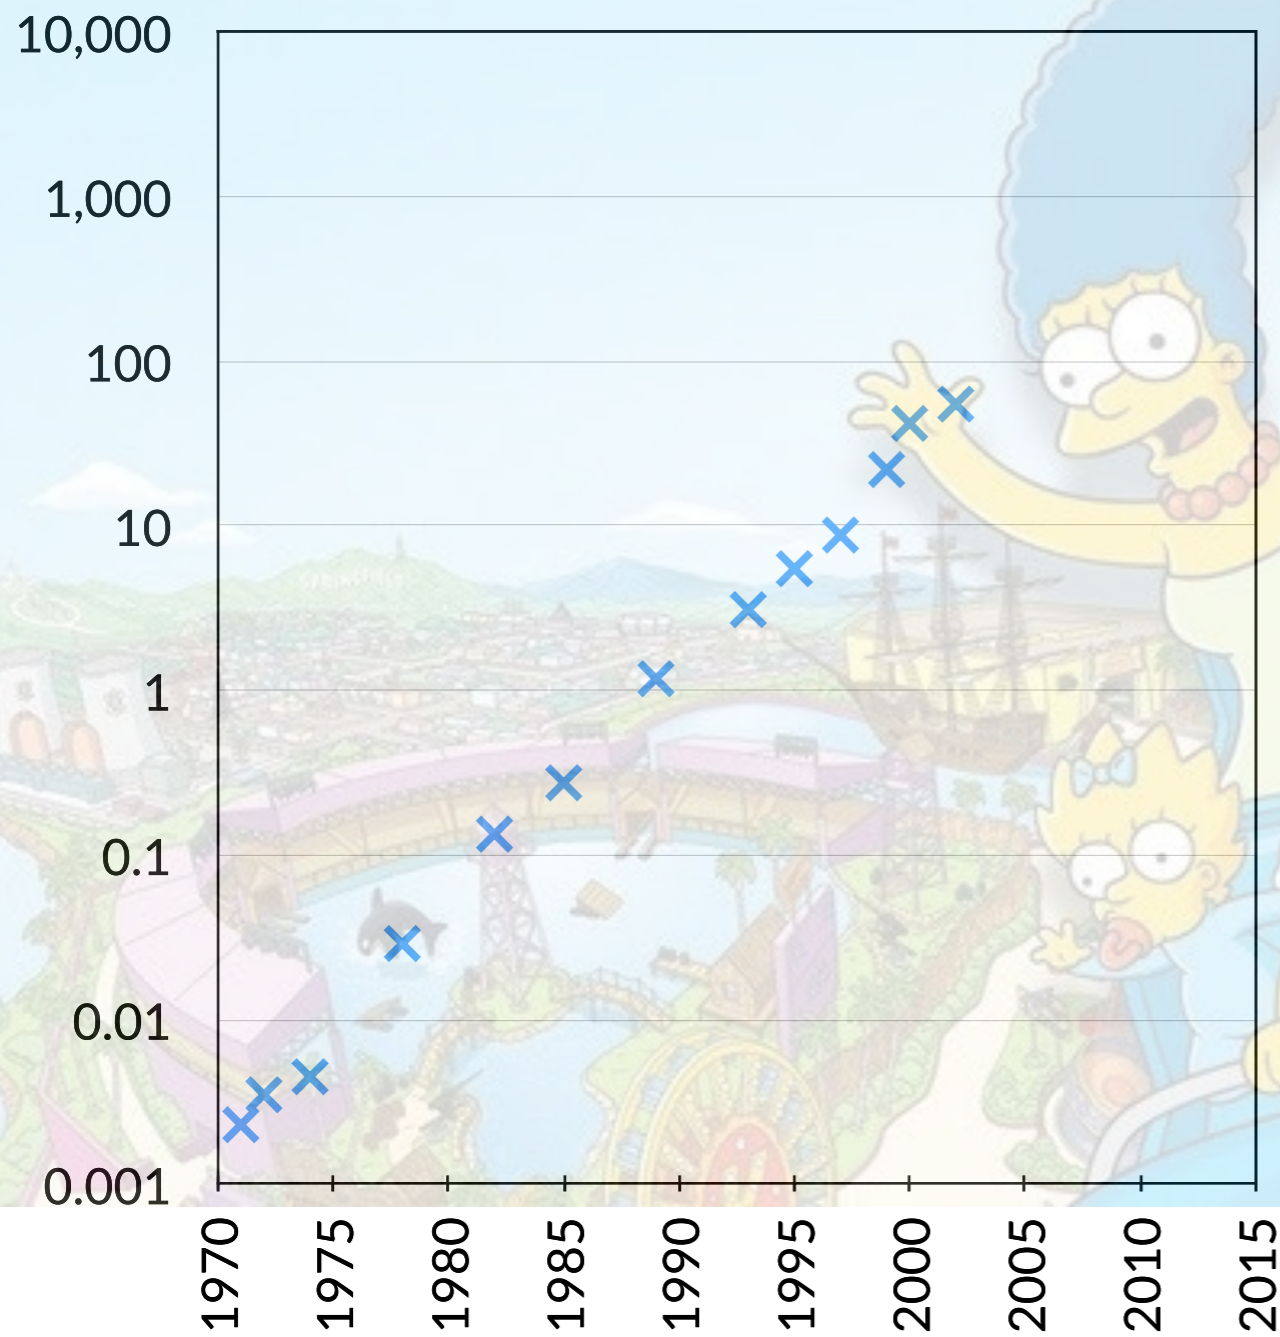
2010

2015

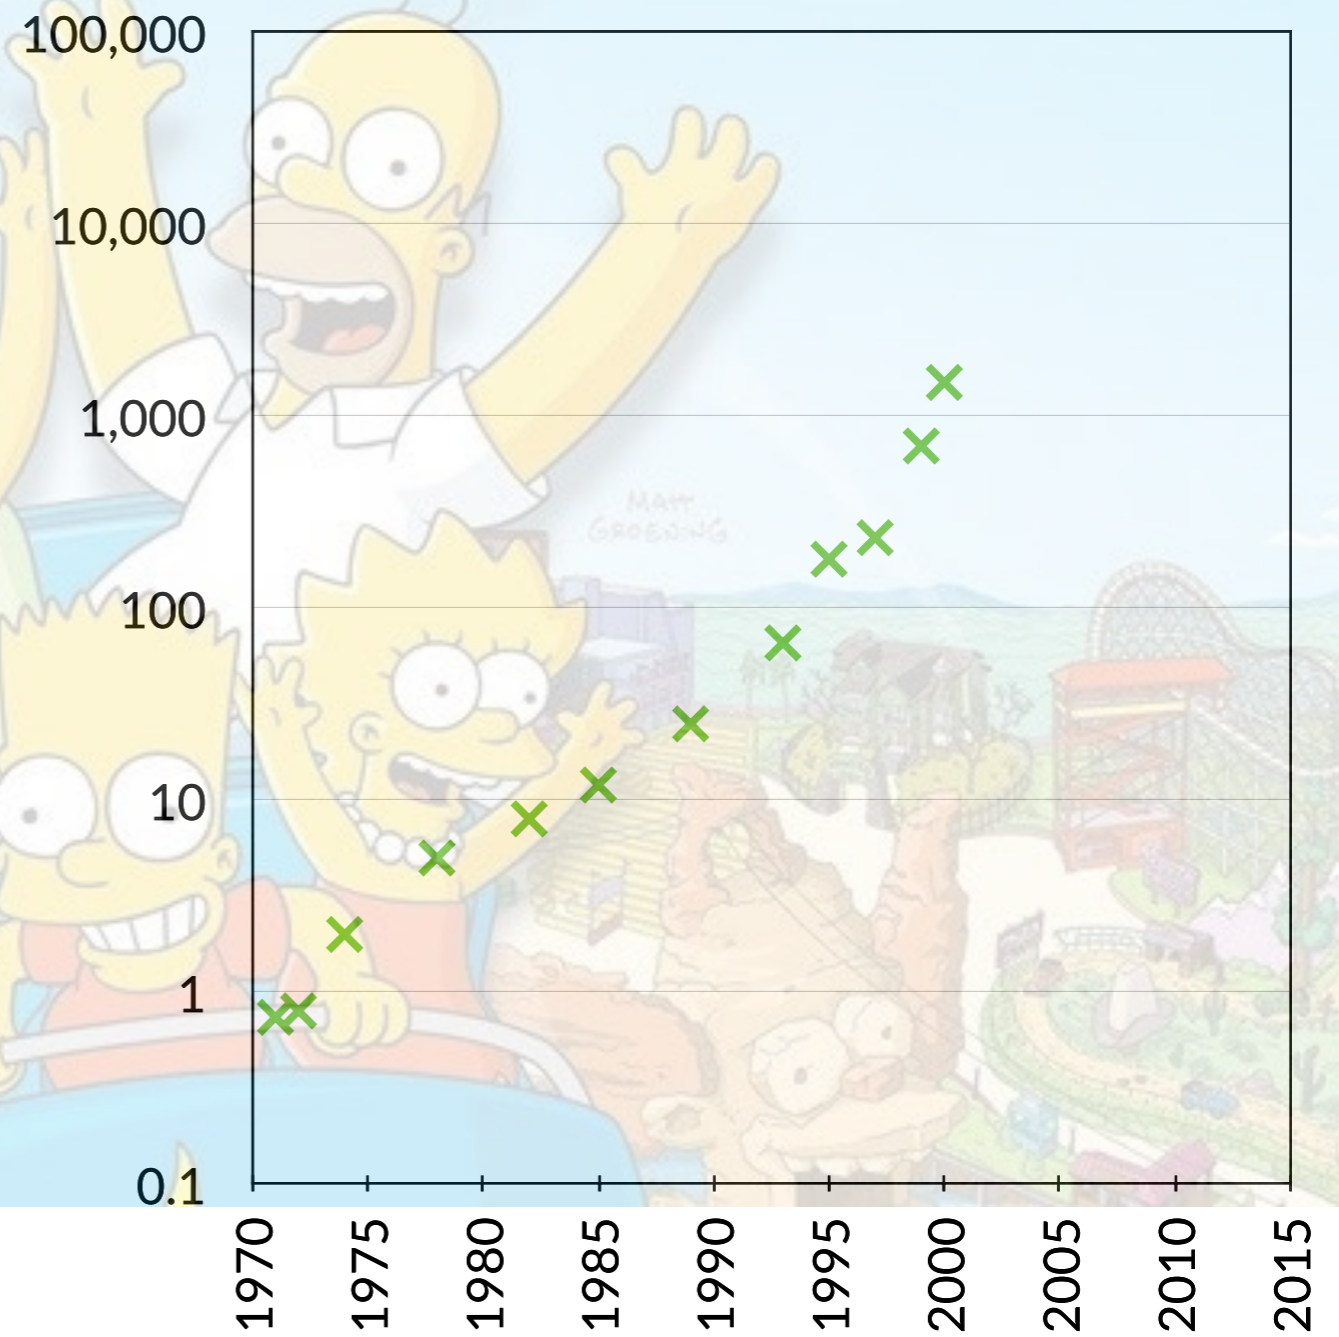
Year

40 Year Performance Ride!

Transistors (millions)



Clock Speed (MHz)

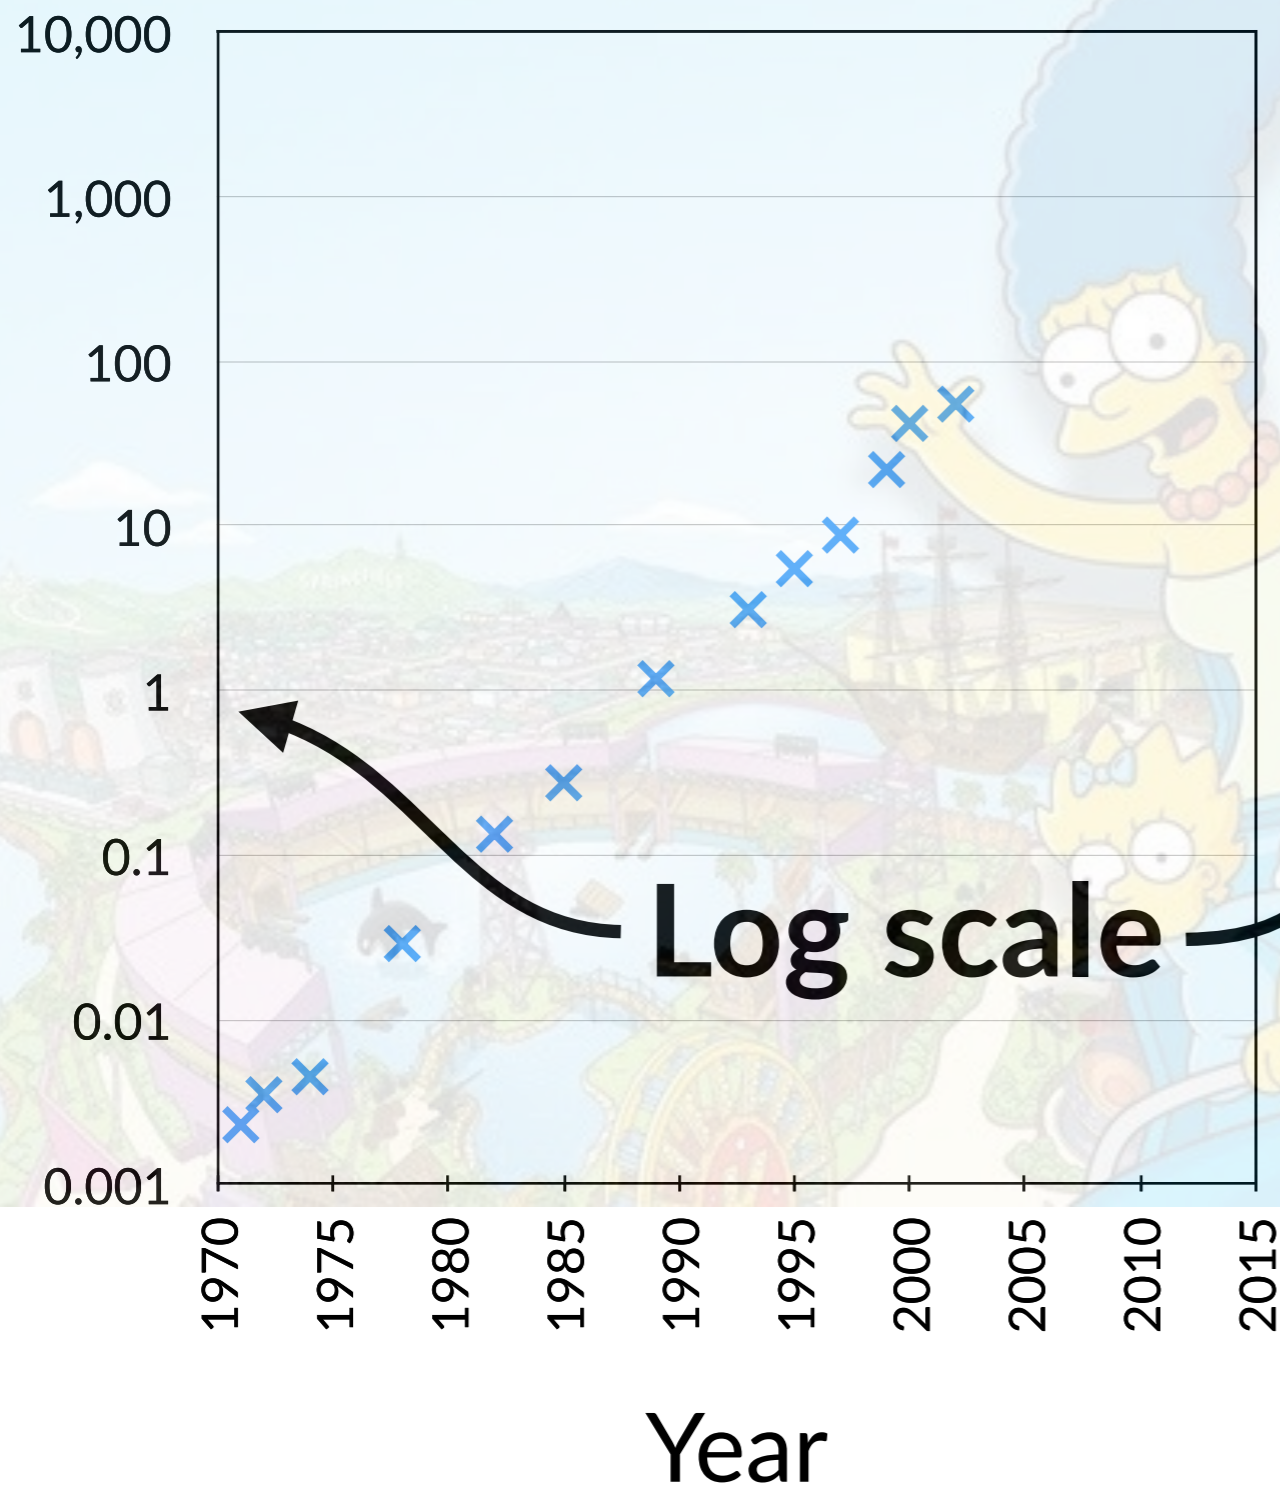


Year

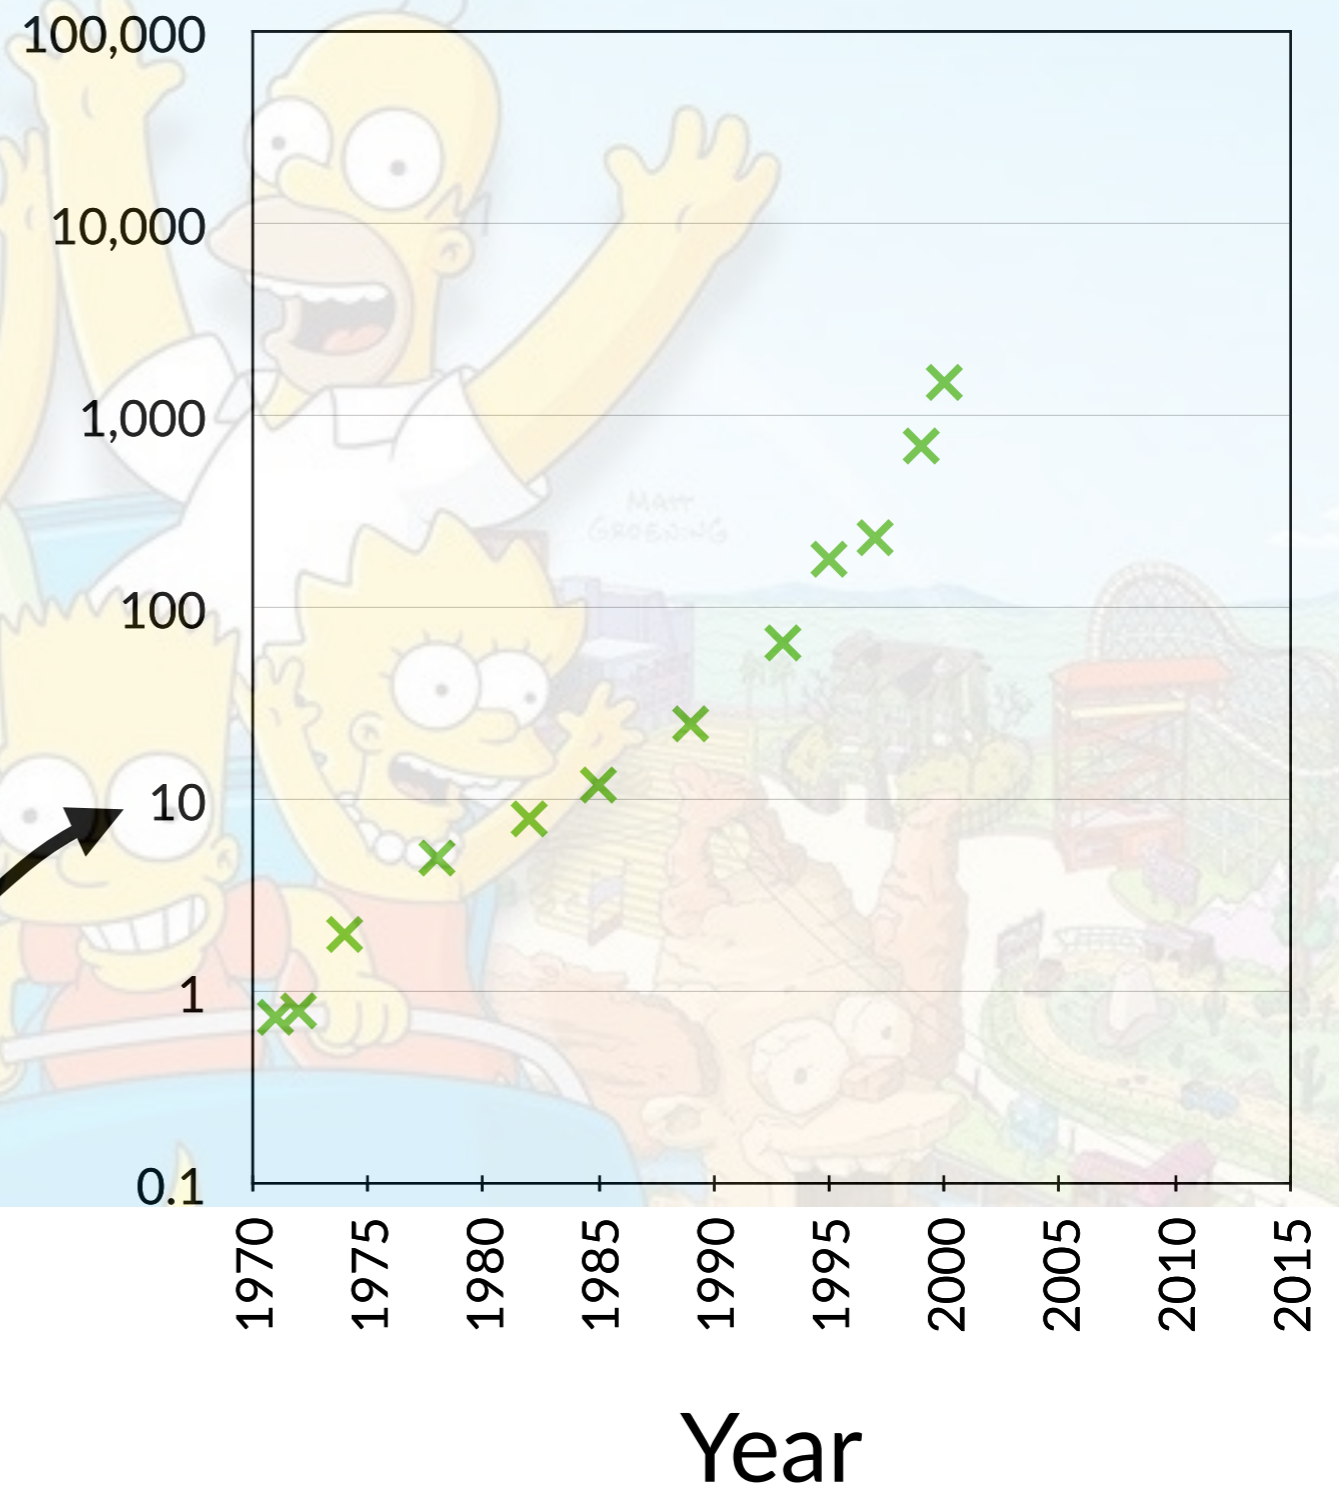
Year

40 Year Performance Ride!

Transistors (millions)

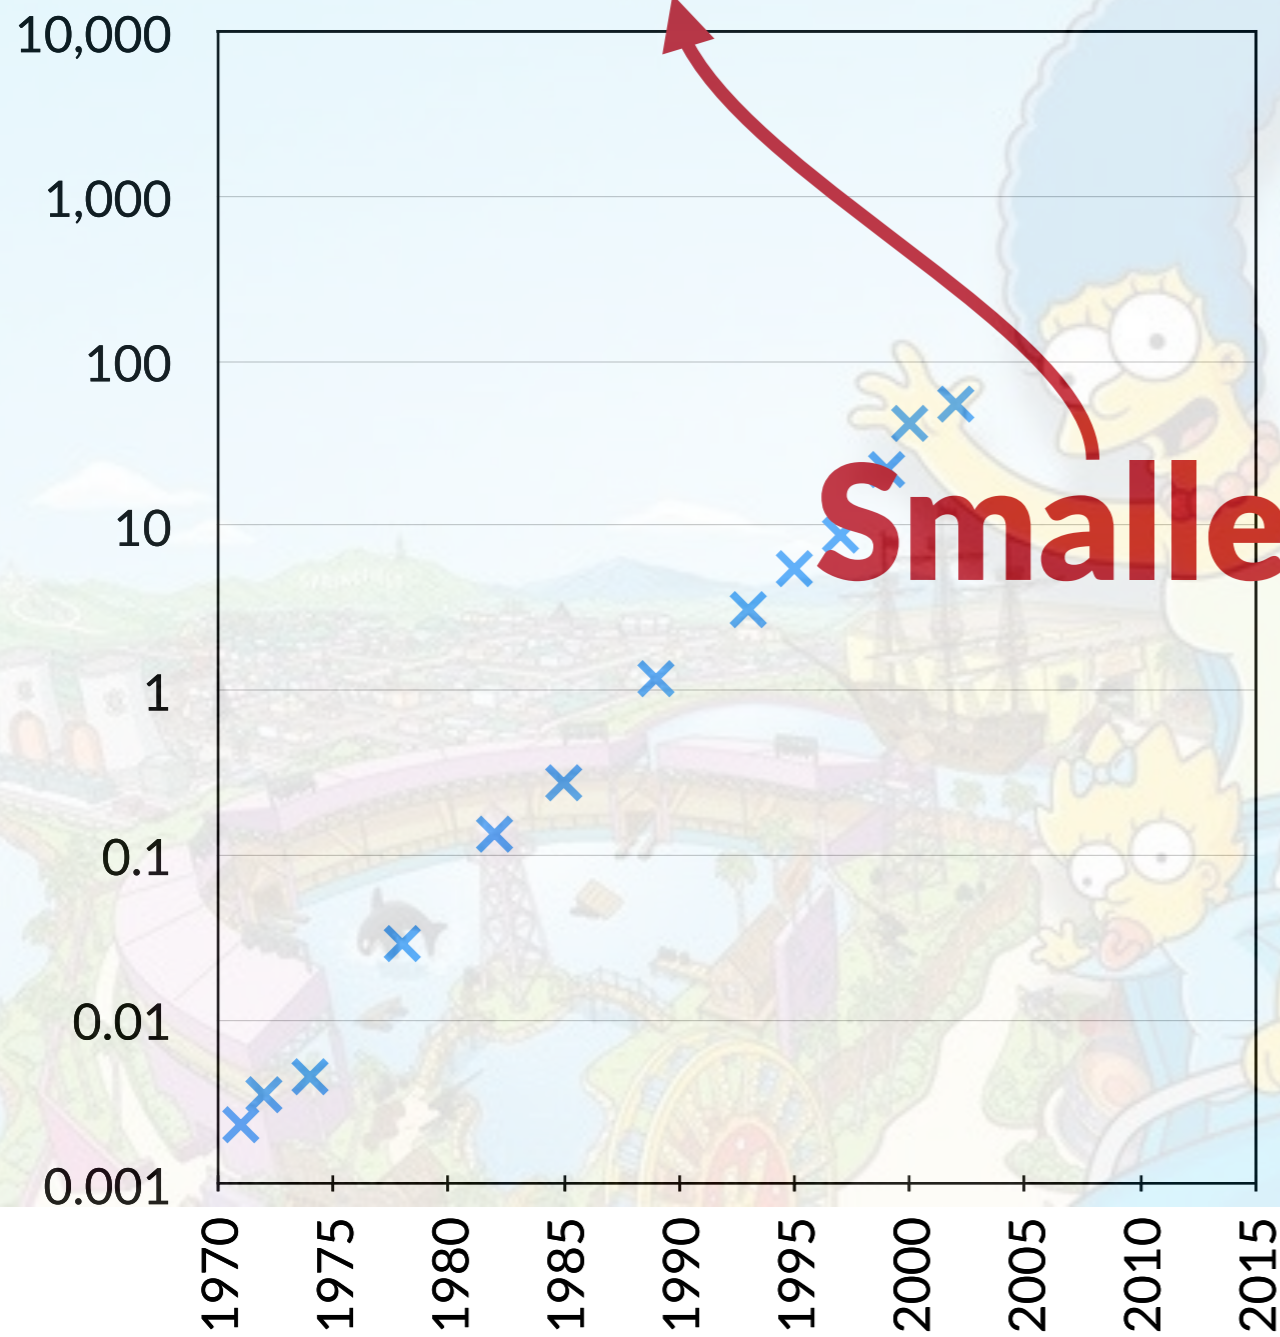


Clock Speed (MHz)

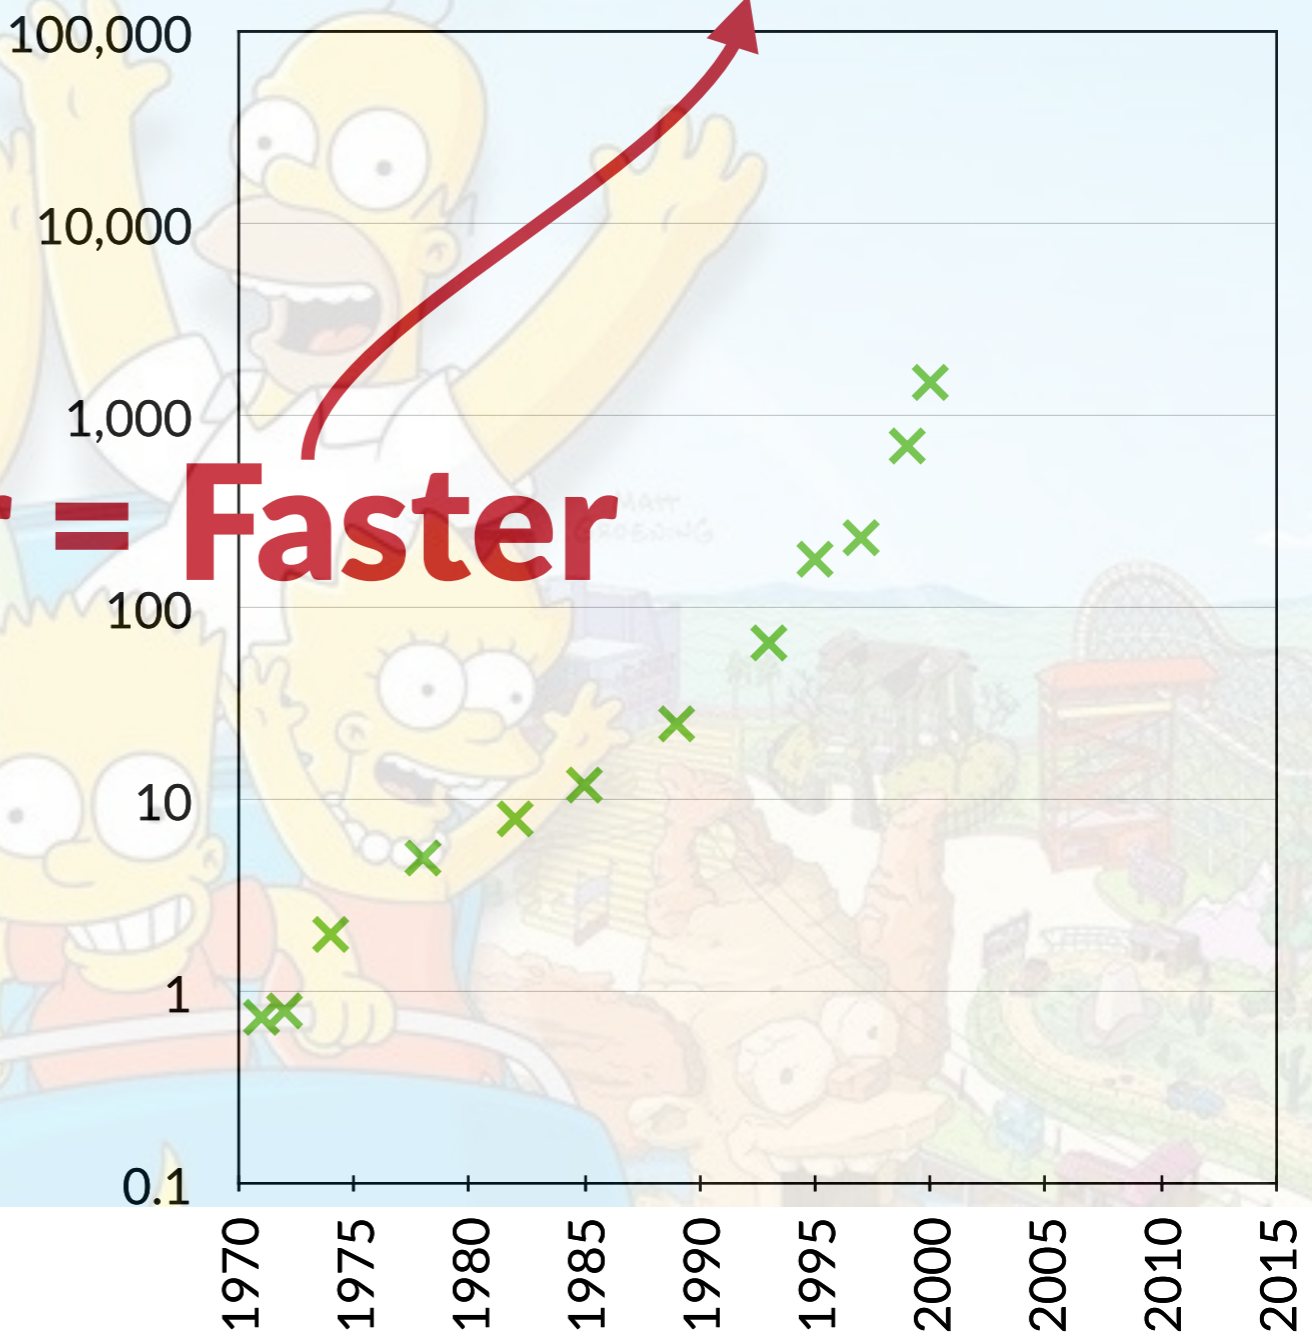


40 Year Performance Ride!

Transistors (millions)



Clock Speed (MHz)



Smaller = Faster

Year


Year

40 Year Performance Ride!

Transistors (millions)




Clock Speed (MHz)






Smaller = Faster

**Just buy new hardware
and your code runs faster**

"Metal" Languages



```
MONITOR FOR IBM 704 1.4 9-14-65 THE MONITOR PAGE 2
CODE 00 00 TO 0000 000 MONITOR
CODE 00 00 TO 0000 000 MONITOR
*****
* FUNCTION: DATA - Initialization
* INPUT: none
* OUTPUT: none
* CALL: none
* DESCRIPTION: none
*****
MONITOR FOR IBM 704 1.4 9-14-65 THE MONITOR PAGE 2
CODE 00 00 TO 0000 000 MONITOR
CODE 00 00 TO 0000 000 MONITOR
*****
* FUNCTION: DATA - Initialization
* INPUT: none
* OUTPUT: none
* CALL: none
* DESCRIPTION: none
*****
MONITOR FOR IBM 704 1.4 9-14-65 THE MONITOR PAGE 2
CODE 00 00 TO 0000 000 MONITOR
CODE 00 00 TO 0000 000 MONITOR
*****
* FUNCTION: DATA - Initialization
* INPUT: none
* OUTPUT: none
* CALL: none
* DESCRIPTION: none
*****
MONITOR FOR IBM 704 1.4 9-14-65 THE MONITOR PAGE 2
CODE 00 00 TO 0000 000 MONITOR
CODE 00 00 TO 0000 000 MONITOR
*****
* FUNCTION: DATA - Initialization
* INPUT: none
* OUTPUT: none
* CALL: none
* DESCRIPTION: none
*****
MONITOR FOR IBM 704 1.4 9-14-65 THE MONITOR PAGE 2
CODE 00 00 TO 0000 000 MONITOR
CODE 00 00 TO 0000 000 MONITOR
*****
* FUNCTION: DATA - Initialization
* INPUT: none
* OUTPUT: none
* CALL: none
* DESCRIPTION: none
*****
```




1949

1957


1972

1985

1980-1990s



1980-1990s



1991




1987



1995




1995



1990

"Irrational Exuberance" Languages



1991



1987


Perl



1995




1995




python™


1990

"Irrational Exuberance" Languages






HOW MANY BYTES IN AN...INT?



HOW MANY BYTES IN AN...INT?






HOW MANY BYTES IN AN...INT?




`sizeof(1)`



HOW MANY BYTES IN AN...INT?



```
sizeof(1)  
→ 4
```




HOW MANY BYTES IN AN...INT?



```
sizeof(1)  
→ 4
```






HOW MANY BYTES IN AN...INT?




```
sizeof(1)  
→ 4
```



```
>>> sys.getsizeof(1)
```


python™



HOW MANY BYTES IN AN...INT?




```
sizeof(1)  
→ 4
```





```
>>> sys.getsizeof(1)  
24
```

python™



HOW MANY BYTES IN AN...EMPTY LIST?






python™

HOW MANY BYTES IN AN...EMPTY LIST?




```
sizeof(list<int>)
```



HOW MANY BYTES IN AN...EMPTY LIST?



```
sizeof(list<int>)  
→ 24
```





HOW MANY BYTES IN AN...EMPTY LIST?



```
sizeof(list<int>)  
→ 24
```






HOW MANY BYTES IN AN...EMPTY LIST?




```
sizeof(list<int>)  
→ 24
```



```
>>> sys.getsizeof([])
```


python™



HOW MANY BYTES IN AN...EMPTY LIST?




```
sizeof(list<int>)  
→ 24
```




python™

```
>>> sys.getsizeof([])  
56
```



HOW MANY BYTES IN AN...EMPTY DICT?






HOW MANY BYTES IN AN...EMPTY DICT?




```
sizeof(map<int,  
int>)
```



HOW MANY BYTES IN AN...EMPTY DICT?



```
sizeof(map<int,  
int>)  
→ 24
```




HOW MANY BYTES IN AN...EMPTY DICT?



```
sizeof(map<int,  
int>)  
→ 24
```






HOW MANY BYTES IN AN...EMPTY DICT?




```
sizeof(map<int,  
int>)  
→ 24
```



python™


```
>>> sys.getsizeof({})
```

HOW MANY BYTES IN AN...EMPTY DICT?




```
sizeof(map<int,  
int>)  
→ 24
```



python™


```
>>> sys.getsizeof({})  
64
```




HOW MANY BYTES IN AN...EMPTY DICT?




```
sizeof(map<int,  
int>)  
→ 24
```





```
>>> sys.getsizeof({})  
64  
python™ 240 (3.6)
```






HOW MANY BYTES IN AN...EMPTY DICT?















python™






MATRIX-
MATRIX
MULTIPLY

```
for i in range(n):  
    for j in range(n):  
        for k in range(n):  
            C[i][j] += A[i][k] * B[k][j]
```











Matt Groening




≈ 2010: THE RIDE IS OVER

The Ride Is Over

Transistors (millions)




Clock Speed (MHz)




The Ride Is Over

Transistors (millions)




Clock Speed (MHz)




The Ride Is Over

Transistors (millions)




Clock Speed (MHz)




The Ride Is Over

Transistors (millions)




Clock Speed (MHz)




The Ride Is Over

Transistors (millions)




Clock Speed (MHz)




The Ride Is Over

Transistors (millions)




Clock Speed (MHz)




The Ride Is Over

Transistors (millions)



Clock Speed (MHz)



Performance in the '80s



just chill and wait for faster hardware!


Modern-day performance

why is this so hard??




Modern-day performance

why is this so hard??



Typical performance evaluation



```
int main(int argc, char **argv) {
    topFrame = (void**)__builtin_frame_address(0);
    setHandler(Trap::TrapSignal, onTrap);
    setHandler(SIGALRM, onTimer);
    setHandler(SIGSEGV, onFault);
    for(Function* f: functions) {
        f->setTrap();
    }
    setTimer(interval);
    int r = stabilizer_main(argc, argv);
    return r;
}

void setTimer(int msec) {
    struct itimerval timer;


    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;
    timer.it_value.tv_usec = 1000 * (msec % 1000);
    timer.it_interval.tv_sec = 0;
    timer.it_interval.tv_usec = 0;

    setitimer(ITIMER_REAL, &timer, 0);
}
```

```
static void flush_icache(void* begin, size_t size) {
    uintptr_t p = (uintptr_t)begin & ~15UL;
    for (size_t i = 0; i < size; i += 16) {
        asm("icbi 0,%0" : : "r"(p));
        p += 16;
    }
    asm("isync");
}
```

```
DataHeapType* getDataHeap() {
    static char buf[sizeof(DataHeapType)];
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;
    return _theDataHeap;
}
```

Typical performance evaluation



```
int main(int argc, char **argv) {
    topFrame = (void**)__builtin_frame_address(0);
    setHandler(Trap::TrapSignal, onTrap);
    setHandler(SIGALRM, onTimer);
    setHandler(SIGSEGV, onFault);
    for(Function* f: functions) {
        f->setTrap();
    }
    setTimer(interval);
    int r = stabilizer_main(argc, argv);
    return r;
}

void setTimer(int msec) {
    struct itimerval timer;


    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;
    timer.it_value.tv_usec = 1000 * (msec % 1000);
    timer.it_interval.tv_sec = 0;
    timer.it_interval.tv_usec = 0;

    setitimer(ITIMER_REAL, &timer, 0);
}
```

```
static void flush_icache(void* begin, size_t size) {
    uintptr_t p = (uintptr_t)begin & ~15UL;
    size_t meaning_of_life=42;
    for (size_t i = 0; i < size; i += 32) {
        asm("icbi 0,%0" : : "r"(p));
        p += 32;
    }
    for (size_t i = 16; i < size; i += 32) {
        asm("icbi 0,%0" : : "r"(p));
        p += 32;
    }
    asm("isync");
}
```

```
DataHeapType* getDataHeap() {
    static char buf[sizeof(DataHeapType)];
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;
    return _theDataHeap;
}
```


Typical performance evaluation



```
int main(int argc, char **argv) {
    topFrame = (void**)__builtin_frame_address(0);
    setHandler(Trap::TrapSignal, onTrap);
    setHandler(SIGALRM, onTimer);
    setHandler(SIGSEGV, onFault);
    for(Function* f: functions) {
        f->setTrap();
    }
    setTimer(interval);
    int r = stabilizer_main(argc, argv);
    return r;
}

void setTimer(int msec) {
    struct itimerval timer;


    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;
    timer.it_value.tv_usec = 1000 * (msec % 1000);
    timer.it_interval.tv_sec = 0;
    timer.it_interval.tv_usec = 0;

    setitimer(ITIMER_REAL, &timer, 0);
}
```

```
static void flush_icache(void* begin, size_t size) {
    uintptr_t p = (uintptr_t)begin & ~15UL;
    size_t meaning_of_life=42;
    for (size_t i = 0; i < size; i += 32) {
        asm("icbi 0,%0" : : "r"(p));
        p += 32;
    }
    for (size_t i = 16; i < size; i += 32) {
        asm("icbi 0,%0" : : "r"(p));
        p += 32;
    }
    asm("isync");
}
```

```
DataHeapType* getDataHeap() {
    static char buf[sizeof(DataHeapType)];
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;
    return _theDataHeap;
}
```

Typical performance evaluation



```
DataHeapType* getDataHeap() {
    static char buf[sizeof(DataHeapType)];
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;
    return _theDataHeap;
}
```

```
int main(int argc, char **argv) {
    topFrame = (void**)_builtin_frame_address(0);
    setHandler(Trap::TrapSignal, onTrap);
    setHandler(SIGALRM, onTimer);
    setHandler(SIGSEGV, onFault);
    for(Function* f: functions) {
        f->setTrap();
    }
    setTimer(interval);
    int r = stabilizer_main(argc, argv);
    return r;
}


void setTimer(int msec) {
    struct itimerval timer;

    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;
    timer.it_value.tv_usec = 1000 * (msec % 1000);
    timer.it_interval.tv_sec = 0;
    timer.it_interval.tv_usec = 0;

    setitimer(ITIMER_REAL, &timer, 0);
}
```

```
static void flush_icache(void* begin, size_t size) {
    uintptr_t p = (uintptr_t)begin & ~15UL;
    size_t meaning_of_life=42;
    for (size_t i = 0; i < size; i += 32) {
        asm("icbi 0,%0" : : "r"(p));
        p += 32;
    }
    for (size_t i = 16; i < size; i += 32) {
        asm("icbi 0,%0" : : "r"(p));
        p += 32;
    }
    asm("isync");
}
```

Typical performance evaluation



```
DataHeapType* getDataHeap() {
    static char buf[sizeof(DataHeapType)];
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;
    return _theDataHeap;
}
```

```
int main(int argc, char **argv) {
    topFrame = (void**)_builtin_frame_address(0);
    setHandler(Trap::TrapSignal, onTrap);
    setHandler(SIGALRM, onTimer);
    setHandler(SIGSEGV, onFault);
    for(Function* f: functions) {
        f->setTrap();
    }
    setTimer(interval);
    int r = stabilizer_main(argc, argv);
    return r;
}
```


```
static void flush_icache(void* begin, size_t size) {
    uintptr_t p = (uintptr_t)begin & ~15UL;
    size_t meaning_of_life=42;
    for (size_t i = 0; i < size; i += 32) {
        asm("icbi 0,%0" : : "r"(p));
        p += 32;
    }
    for (size_t i = 16; i < size; i += 32) {
        asm("icbi 0,%0" : : "r"(p));
        p += 32;
    }
    asm("isync");
}
```

```
void setTimer(int msec) {
    struct itimerval timer;

    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;
    timer.it_value.tv_usec = 1000 * (msec % 1000);
    timer.it_interval.tv_sec = 0;
    timer.it_interval.tv_usec = 0;

    setitimer(ITIMER_REAL, &timer, 0);
}
```

Typical performance evaluation



```
DataHeapType* getDataHeap() {
    static char buf[sizeof(DataHeapType)];
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;
    return _theDataHeap;
}
```

```
int main(int argc, char **argv) {
    topFrame = (void**)_builtin_frame_address(0);
    setHandler(Trap::TrapSignal, onTrap);
    setHandler(SIGALRM, onTimer);
    setHandler(SIGSEGV, onFault);
    for(Function* f: functions) {
        f->setTrap();
    }
    setTimer(interval);
    int r = stabilizer_main(argc, argv);
    return r;
}
```

```
static void flush_icache(void* begin, size_t size) {
    uintptr_t p = (uintptr_t)begin & ~15UL;
    size_t meaning_of_life=42;
    for (size_t i = 0; i < size; i += 32) {
        asm("icbi 0,%0" : : "r"(p));
        p += 32;
    }
    for (size_t i = 16; i < size; i += 32) {
        asm("icbi 0,%0" : : "r"(p));
        p += 32;
    }
    asm("isync");
}
```

```
void setTimer(int msec) {
    struct itimerval timer;

    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;
    timer.it_value.tv_usec = 1000 * (msec % 1000);
    timer.it_interval.tv_sec = 0;
    timer.it_interval.tv_usec = 0;

    setitimer(ITIMER_REAL, &timer, 0);
}
```

Typical performance evaluation

A

A'

```
DataHeapType* getDataHeap() {
    static char buf[sizeof(DataHeapType)];
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;
    return _theDataHeap;
}
```

```
int main(int argc, char **argv) {
    topFrame = (void**)_builtin_frame_address(0);
    setHandler(Trap::TrapSignal, onTrap);
    setHandler(SIGALRM, onTimer);
    setHandler(SIGSEGV, onFault);
    for(Function* f: functions) {
        f->setTrap();
    }
    setTimer(interval);
    int r = stabilizer_main(argc, argv);
    return r;
}
```

```
static void flush_icache(void* begin, size_t size) {
    uintptr_t p = (uintptr_t)begin & ~15UL;
    size_t meaning_of_life=42;
    for (size_t i = 0; i < size; i += 32) {
        asm("icbi 0,%0" : "r"(p));
        p += 32;
    }
    for (size_t i = 16; i < size; i += 32) {
        asm("icbi 0,%0" : "r"(p));
        p += 32;
    }
    asm("isync");
}
```



```
void setTimer(int msec) {
    struct itimerval timer;

    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;
    timer.it_value.tv_usec = 1000 * (msec % 1000);
    timer.it_interval.tv_sec = 0;
    timer.it_interval.tv_usec = 0;


    setitimer(ITIMER_REAL, &timer, 0);
}
```




Which is faster?




Which is faster?




90s



Which is faster?




90s




87.5s


Which is faster?



90s



87.5s



2.8% faster!



Why is  faster than  ?

Why is faster than ?

```
DataHeapType* getDataHeap() {
    static char buf[sizeof(DataHeapType)];
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;
    return _theDataHeap;
}
```

```
int main(int argc, char **argv) {
    topFrame = (void**)__builtin_frame_address(0);
    setHandler(Trap::TrapSignal, onTrap);
    setHandler(SIGALRM, onTimer);
    setHandler(SIGSEGV, onFault);
    for(Function* f: functions) {
        f->setTrap();
    }
    setTimer(interval);
    int r = stabilizer_main(argc, argv);
    return r;
}
```

```
static void flush_icache(void* begin, size_t size) {
    uintptr_t p = (uintptr_t)begin & ~15UL;
    for (size_t i = 0; i < size; i += 16) {
        asm("icbi 0,%0" : : "r"(p));
        p += 16;
    }
}
```

```
void setTimer(int msec) {
    struct itimerval timer;

    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;
    timer.it_value.tv_usec = 1000 * (msec % 1000);
    timer.it_interval.tv_sec = 0;
    timer.it_interval.tv_usec = 0;

    setitimer(ITIMER_REAL, &timer, 0);
}
```

```
int main(int argc, char **argv) {
    topFrame = (void**)__builtin_frame_address(0);
    setHandler(Trap::TrapSignal, onTrap);
    setHandler(SIGALRM, onTimer);
    setHandler(SIGSEGV, onFault);
    for(Function* f: functions) {
        f->setTrap();
    }
    setTimer(interval);
    int r = stabilizer_main(argc, argv);
    return r;
}

void setTimer(int msec) {
    struct itimerval timer;

    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;
    timer.it_value.tv_usec = 1000 * (msec % 1000);
    timer.it_interval.tv_sec = 0;
    timer.it_interval.tv_usec = 0;

    setitimer(ITIMER_REAL, &timer, 0);
}
```

```
static void flush_icache(void* begin, size_t size) {
    uintptr_t p = (uintptr_t)begin & ~15UL;
    for (size_t i = 0; i < size; i += 16) {
        asm("icbi 0,%0" : : "r"(p));
        p += 16;
    }
    asm("isync");
}
```

```
DataHeapType* getDataHeap() {
    static char buf[sizeof(DataHeapType)];
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;
    return _theDataHeap;
}
```

Why is A' faster than A ?

```
DataHeapType* getDataHeap() {
    static char buf[sizeof(DataHeapType)];
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;
    return _theDataHeap;
}
```

```
int main(int argc, char **argv) {
    topFrame = (void**)__builtin_frame_address(0);
    setHandler(Trap::TrapSignal, onTrap);
    setHandler(SIGALRM, onTimer);
    setHandler(SIGSEGV, onFault);
    for(Function* f: functions) {
        f->setTrap();
    }
    setTimer(interval);
    int r = stabilizer_main(argc, argv);
    return r;
}
```

```
static void flush_icache(void* begin, size_t size) {
    uintptr_t p = (uintptr_t)begin & ~15UL;
    for (size_t i = 0; i < size; i += 16) {
        asm("icbi 0,%0" : : "r"(p));
        p += 16;
    }
}
```

```
void setTimer(int msec) {
    struct itimerval timer;

    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;
    timer.it_value.tv_usec = 1000 * (msec % 1000);
    timer.it_interval.tv_sec = 0;
    timer.it_interval.tv_usec = 0;

    setitimer(ITIMER_REAL, &timer, 0);
}
```

```
int main(int argc, char **argv) {
    topFrame = (void**)__builtin_frame_address(0);
    setHandler(Trap::TrapSignal, onTrap);
    setHandler(SIGALRM, onTimer);
    setHandler(SIGSEGV, onFault);
    for(Function* f: functions) {
        f->setTrap();
    }
    setTimer(interval);
    int r = stabilizer_main(argc, argv);
    return r;
}

void setTimer(int msec) {
    struct itimerval timer;

    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;
    timer.it_value.tv_usec = 1000 * (msec % 1000);
    timer.it_interval.tv_sec = 0;
    timer.it_interval.tv_usec = 0;

    setitimer(ITIMER_REAL, &timer, 0);
}
```

```
static void flush_icache(void* begin, size_t size) {
    uintptr_t p = (uintptr_t)begin & ~15UL;
    for (size_t i = 0; i < size; i += 16) {
        asm("icbi 0,%0" : : "r"(p));
        p += 16;
    }
    asm("isync");
}
```

```
DataHeapType* getDataHeap() {
    static char buf[sizeof(DataHeapType)];
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;
    return _theDataHeap;
}
```

The code change!



Why is A' faster than A ?

```
DataHeapType* getDataHeap() {
    static char buf[sizeof(DataHeapType)];
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;
    return _theDataHeap;
}
```

```
int main(int argc, char **argv) {
    topFrame = (void**)__builtin_frame_address(0);
    setHandler(Trap::TrapSignal, onTrap);
    setHandler(SIGALRM, onTimer);
    setHandler(SIGSEGV, onFault);
    for(Function* f: functions) {
        f->setTrap();
    }
    setTimer(interval);
    int r = stabilizer_main(argc, argv);
    return r;
}
```

```
static void flush_icache(void* begin, size_t size) {
    uintptr_t p = (uintptr_t)begin & ~15UL;
    for (size_t i = 0; i < size; i += 16) {
        asm("icbi 0,%0" : : "r"(p));
        p += 16;
    }
}
```

```
void setTimer(int msec) {
    struct itimerval timer;

    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;
    timer.it_value.tv_usec = 1000 * (msec % 1000);
    timer.it_interval.tv_sec = 0;
    timer.it_interval.tv_usec = 0;

    setitimer(ITIMER_REAL, &timer, 0);
}
```

```
int main(int argc, char **argv) {
    topFrame = (void**)__builtin_frame_address(0);
    setHandler(Trap::TrapSignal, onTrap);
    setHandler(SIGALRM, onTimer);
    setHandler(SIGSEGV, onFault);
    for(Function* f: functions) {
        f->setTrap();
    }
    setTimer(interval);
    int r = stabilizer_main(argc, argv);
    return r;
}

void setTimer(int msec) {
    struct itimerval timer;


    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;
    timer.it_value.tv_usec = 1000 * (msec % 1000);
    timer.it_interval.tv_sec = 0;
    timer.it_interval.tv_usec = 0;

    setitimer(ITIMER_REAL, &timer, 0);
}
```

```
static void flush_icache(void* begin, size_t size) {
    uintptr_t p = (uintptr_t)begin & ~15UL;
    for (size_t i = 0; i < size; i += 16) {
        asm("icbi 0,%0" : : "r"(p));
        p += 16;
    }
    asm("isync");
}
```

```
DataHeapType* getDataHeap() {
    static char buf[sizeof(DataHeapType)];
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;
    return _theDataHeap;
}
```

The code change!



Why is **A'** faster than **A** ?

Or was it
an accident?

```
DataHeapType* getDataHeap() {  
    static char buf[sizeof(DataHeapType)];  
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;  
    return _theDataHeap;  
}
```

```
int main(int argc, char **argv) {  
    topFrame = (void**)__builtin_frame_address(0);  
    setHandler(Trap::TrapSignal, onTrap);  
    setHandler(SIGALRM, onTimer);  
    setHandler(SIGSEGV, onFault);  
    for(Function* f: functions) {  
        f->setTrap();  
    }  
    setTimer(interval);  
    int r = stabilizer_main(argc, argv);  
    return r;  
}
```

```
static void flush_icache(void* begin, size_t size) {  
    uintptr_t p = (uintptr_t)begin & ~15UL;  
    for (size_t i = 0; i < size; i += 16) {  
        asm("icbi 0,%0" : : "r"(p));  
        p += 16;  
    }  
}
```

```
void setTimer(int msec) {  
    struct itimerval timer;  
  
    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;  
    timer.it_value.tv_usec = 1000 * (msec % 1000);  
    timer.it_interval.tv_sec = 0;  
    timer.it_interval.tv_usec = 0;  
  
    setitimer(ITIMER_REAL, &timer, 0);  
}
```

```
int main(int argc, char **argv) {  
    topFrame = (void**)__builtin_frame_address(0);  
    setHandler(Trap::TrapSignal, onTrap);  
    setHandler(SIGALRM, onTimer);  
    setHandler(SIGSEGV, onFault);  
    for(Function* f: functions) {  
        f->setTrap();  
    }  
    setTimer(interval);  
    int r = stabilizer_main(argc, argv);  
    return r;  
}  
  
void setTimer(int msec) {  
    struct itimerval timer;  
  
    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;  
    timer.it_value.tv_usec = 1000 * (msec % 1000);  
    timer.it_interval.tv_sec = 0;  
    timer.it_interval.tv_usec = 0;  
  
    setitimer(ITIMER_REAL, &timer, 0);  
}
```

```
static void flush_icache(void* begin, size_t size) {  
    uintptr_t p = (uintptr_t)begin & ~15UL;  
    for (size_t i = 0; i < size; i += 16) {  
        asm("icbi 0,%0" : : "r"(p));  
        p += 16;  
    }  
    asm("isync");  
}
```

```
DataHeapType* getDataHeap() {  
    static char buf[sizeof(DataHeapType)];  
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;  
    return _theDataHeap;  
}
```



Why is A' faster than A ?

Or was it
an accident?

```
DataHeapType* getDataHeap() {  
    static char buf[sizeof(DataHeapType)];  
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;  
    return _theDataHeap;  
}
```

```
int main(int argc, char **argv) {  
    topFrame = (void**)__builtin_frame_address(0);  
    setHandler(Trap::TrapSignal, onTrap);  
    setHandler(SIGALRM, onTimer);  
    setHandler(SIGSEGV, onFault);  
    for(Function* f: functions) {  
        f->setTrap();  
    }  
    setTimer(interval);  
    int r = stabilizer_main(argc, argv);  
    return r;  
}
```

```
static void flush_icache(void* begin, size_t size) {  
    uintptr_t p = (uintptr_t)begin & ~15UL;  
    for (size_t i = 0; i < size; i += 16) {  
        asm("icbi 0,%0" : : "r"(p));  
        p += 16;  
    }  
}
```

```
void setTimer(int msec) {  
    struct itimerval timer;  
  
    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;  
    timer.it_value.tv_usec = 1000 * (msec % 1000);  
    timer.it_interval.tv_sec = 0;  
    timer.it_interval.tv_usec = 0;  
  
    setitimer(ITIMER_REAL, &timer, 0);  
}
```

```
int main(int argc, char **argv) {  
    topFrame = (void**)__builtin_frame_address(0);  
    setHandler(Trap::TrapSignal, onTrap);  
    setHandler(SIGALRM, onTimer);  
    setHandler(SIGSEGV, onFault);  
    for(Function* f: functions) {  
        f->setTrap();  
    }  
    setTimer(interval);  
    int r = stabilizer_main(argc, argv);  
    return r;  
}  
  
void setTimer(int msec) {  
    struct itimerval timer;  
  
    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;  
    timer.it_value.tv_usec = 1000 * (msec % 1000);  
    timer.it_interval.tv_sec = 0;  
    timer.it_interval.tv_usec = 0;  
  
    setitimer(ITIMER_REAL, &timer, 0);  
}
```

```
static void flush_icache(void* begin, size_t size) {  
    uintptr_t p = (uintptr_t)begin & ~15UL;  
    for (size_t i = 0; i < size; i += 16) {  
        asm("icbi 0,%0" : : "r"(p));  
        p += 16;  
    }  
    asm("isync");  
}
```

```
DataHeapType* getDataHeap() {  
    static char buf[sizeof(DataHeapType)];  
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;  
    return _theDataHeap;  
}
```



Why is **A'** faster than **A** ?

Layout biases measurement

Mytkowicz et al. (ASPLOS'09)

```
DataHeapType* getDataHeap() {  
    static char buf[sizeof(DataHeapType)];  
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;  
    return _theDataHeap;  
}
```

```
int main(int argc, char **argv) {  
    topFrame = (void**)_builtin_frame_address(0);  
    setHandler(Trap::TrapSignal, onTrap);  
    setHandler(SIGALRM, onTimer);  
    setHandler(SIGSEGV, onFault);  
    for(Function* f: functions) {  
        f->setTrap();  
    }  
    setTimer(interval);  
    int r = stabilizer_main(argc, argv);  
    return r;  
}
```

```
static void flush_icache(void* begin, size_t size) {  
    uintptr_t p = (uintptr_t)begin & ~15UL;  
    for (size_t i = 0; i < size; i += 16) {  
        asm("icbi 0,%0" : : "r"(p));  
        p += 16;  
    }  
}
```

```
void setTimer(int msec) {  
    struct itimerval timer;  
  
    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;  
    timer.it_value.tv_usec = 1000 * (msec % 1000);  
    timer.it_interval.tv_sec = 0;  
    timer.it_interval.tv_usec = 0;  
  
    setitimer(ITIMER_REAL, &timer, 0);  
}
```

```
int main(int argc, char **argv) {  
    topFrame = (void**)_builtin_frame_address(0);  
    setHandler(Trap::TrapSignal, onTrap);  
    setHandler(SIGALRM, onTimer);  
    setHandler(SIGSEGV, onFault);  
    for(Function* f: functions) {  
        f->setTrap();  
    }  
    setTimer(interval);  
    int r = stabilizer_main(argc, argv);  
    return r;  
}  
  
void setTimer(int msec) {  
    struct itimerval timer;  
  
    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;  
    timer.it_value.tv_usec = 1000 * (msec % 1000);  
    timer.it_interval.tv_sec = 0;  
    timer.it_interval.tv_usec = 0;  
  
    setitimer(ITIMER_REAL, &timer, 0);  
}
```

```
static void flush_icache(void* begin, size_t size) {  
    uintptr_t p = (uintptr_t)begin & ~15UL;  
    for (size_t i = 0; i < size; i += 16) {  
        asm("icbi 0,%0" : : "r"(p));  
        p += 16;  
    }  
    asm("isync");  
}
```

```
DataHeapType* getDataHeap() {  
    static char buf[sizeof(DataHeapType)];  
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;  
    return _theDataHeap;  
}
```



Layout biases measurement

Mytkowicz et al. (ASPLOS'09)

Producing Wrong Data Without Doing Anything Obviously Wrong!

Todd Mytkowicz Amer Diwan

Department of Computer Science
University of Colorado
Boulder, CO, USA

{mytkowit,diwan}@colorado.edu

Matthias Hauswirth

Faculty of Informatics
University of Lugano
Lugano, CH

Matthias.Hauswirth@unisi.ch

Peter F. Sweeney

IBM Research
Hawthorne, NY, USA

pfs@us.ibm.com

Layout biases measurement

Mytkowicz et al. (ASPLOS'09)

Producing Wrong Data Without Doing Anything Obviously Wrong!

Todd Mytkowicz Amer Diwan

Department of Computer Science
University of Colorado
Boulder, CO, USA

{mytkowit,diwan}@colorado.edu

Matthias Hauswirth

Faculty of Informatics
University of Lugano
Lugano, CH

Matthias.Hauswirth@unisi.ch

Peter F. Sweeney

IBM Research
Hawthorne, NY, USA

pfs@us.ibm.com

Link Order

Changes function addresses

Layout biases measurement

Mytkowicz et al. (ASPLOS'09)

Producing Wrong Data Without Doing Anything Obviously Wrong!

Todd Mytkowicz Amer Diwan

Department of Computer Science
University of Colorado
Boulder, CO, USA

{mytkowit,diwan}@colorado.edu

Matthias Hauswirth

Faculty of Informatics
University of Lugano
Lugano, CH

Matthias.Hauswirth@unisi.ch

Peter F. Sweeney

IBM Research
Hawthorne, NY, USA

pfs@us.ibm.com

Link Order

Changes function addresses

Environment

Variable Size

Moves the program stack

Layout biases measurement

Mytkowicz et al. (ASPLOS'09)

Link Order

Changes function addresses

**Larger than the impact
of -O3! ($\pm 40\%$)**

Environment

Variable Size

Moves the program stack

Why is faster than ?



```
int main(int argc, char **argv) {
    topFrame = (void**)__builtin_frame_address(0);
    setHandler(Trap::TrapSignal, onTrap);
    setHandler(SIGALRM, onTimer);
    setHandler(SIGSEGV, onFault);
    for(Function* f: functions) {
        f->setTrap();
    }
    setTimer(interval);
    int r = stabilizer_main(argc, argv);
    return r;
}

void setTimer(int msec) {
    struct itimerval timer;




    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;
    timer.it_value.tv_usec = 1000 * (msec % 1000);
    timer.it_interval.tv_sec = 0;
    timer.it_interval.tv_usec = 0;

    setitimer(ITIMER_REAL, &timer, 0);
}
```

```
static void flush_icache(void* begin, size_t size) {
    uintptr_t p = (uintptr_t)begin & ~15UL;
    for (size_t i = 0; i < size; i += 16) {
        asm("icbi 0,%0" : : "r"(p));
        p += 16;
    }
    asm("isync");
}
```

```
DataHeapType* getDataHeap() {
    static char buf[sizeof(DataHeapType)];
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;
    return _theDataHeap;
}
```

Why is **A'** faster than **A** ?



```
int main(int argc, char **argv) {
    topFrame = (void**)__builtin_frame_address(0);
    setHandler(Trap::TrapSignal, onTrap);
    setHandler(SIGALRM, onTimer);
    setHandler(SIGSEGV, onFault);
    for(Function* f: functions) {
        f->setTrap();
    }
    setTimer(interval);
    int r = stabilizer_main(argc, argv);
    return r;
}

void setTimer(int msec) {
    struct itimerval timer;

    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;
    timer.it_value.tv_usec = 1000 * (msec % 1000);
    timer.it_interval.tv_sec = 0;
    timer.it_interval.tv_usec = 0;

    setitimer(ITIMER_REAL, &timer, 0);
}

```

```
static void flush_icache(void* begin, size_t size) {
    uintptr_t p = (uintptr_t)begin & ~15UL;
    for (size_t i = 0; i < size; i += 16) {
        asm("icbi 0,%0" : : "r"(p));
        p += 16;
    }
    asm("isync");
}

```



```
DataHeapType* getDataHeap() {
    static char buf[sizeof(DataHeapType)];
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;
    return _theDataHeap;
}

```



map to same
cache set

Why is **A'** faster than **A** ?



```
int main(int argc, char **argv) {
    topFrame = (void**)__builtin_frame_address(0);
    setHandler(Trap::TrapSignal, onTrap);
    setHandler(SIGALRM, onTimer);
    setHandler(SIGSEGV, onFault);
    for(Function* f: functions) {
        f->setTrap();
    }
    setTimer(interval);
    int r = stabilizer_main(argc, argv);
    return r;
}

void setTimer(int msec) {
    struct itimerval timer;

    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;
    timer.it_value.tv_usec = 1000 * (msec % 1000);
    timer.it_interval.tv_sec = 0;
    timer.it_interval.tv_usec = 0;

    setitimer(ITIMER_REAL, &timer, 0);
}
```

```
static void flush_icache(void* begin, size_t size) {
    uintptr_t p = (uintptr_t)begin & ~15UL;
    for (size_t i = 0; i < size; i += 16) {
        asm("icbi 0,%0" : : "r"(p));
        p += 16;
    }
    asm("isync");
}
```

```
DataHeapType* getDataHeap() {
    static char buf[sizeof(DataHeapType)];
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;
    return _theDataHeap;
}
```



map to same
cache set
conflict

Why is **A'** faster than **A** ?



```
int main(int argc, char **argv) {
    topFrame = (void**)__builtin_frame_address(0);
    setHandler(Trap::TrapSignal, onTrap);
    setHandler(SIGALRM, onTimer);
    setHandler(SIGSEGV, onFault);
    for(Function* f: functions) {
        f->setTrap();
    }
    setTimer(interval);
    int r = stabilizer_main(argc, argv);
    return r;
}


void setTimer(int msec) {
    struct itimerval timer;

    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;
    timer.it_value.tv_usec = 1000 * (msec % 1000);
    timer.it_interval.tv_sec = 0;
    timer.it_interval.tv_usec = 0;

    setitimer(ITIMER_REAL, &timer, 0);
}
```

```
static void flush_icache(void* begin, size_t size) {
    uintptr_t p = (uintptr_t)begin & ~15UL;
    for (size_t i = 0; i < size; i += 16) {
        asm("icbi 0,%0" : : "r"(p));
        p += 16;
    }
    asm("isync");
}
```

```
DataHeapType* getDataHeap() {
    static char buf[sizeof(DataHeapType)];
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;
    return _theDataHeap;
}
```



```
DataHeapType* getDataHeap() {
    static char buf[sizeof(DataHeapType)];
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;
    return _theDataHeap;
}
```

```
int main(int argc, char **argv) {
    topFrame = (void**)__builtin_frame_address(0);
    setHandler(Trap::TrapSignal, onTrap);
    setHandler(SIGALRM, onTimer);
    setHandler(SIGSEGV, onFault);
    for(Function* f: functions) {
        f->setTrap();
    }
    setTimer(interval);
    int r = stabilizer_main(argc, argv);
    return r;
}
```



```
static void flush_icache(void* begin, size_t size) {
    uintptr_t p = (uintptr_t)begin & ~15UL;
    for (size_t i = 0; i < size; i += 16) {
        asm("icbi 0,%0" : : "r"(p));
        p += 16;
    }
}
```

```
void setTimer(int msec) {
    struct itimerval timer;

    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;
    timer.it_value.tv_usec = 1000 * (msec % 1000);
    timer.it_interval.tv_sec = 0;
    timer.it_interval.tv_usec = 0;

    setitimer(ITIMER_REAL, &timer, 0);
}
```


Why is **A'** faster than **A** ?



```
int main(int argc, char **argv) {
    topFrame = (void**)__builtin_frame_address(0);
    setHandler(Trap::TrapSignal, onTrap);
    setHandler(SIGALRM, onTimer);
    setHandler(SIGSEGV, onFault);
    for(Function* f: functions) {
        f->setTrap();
    }
    setTimer(interval);
    int r = stabilizer_main(argc, argv);
    return r;
}

void setTimer(int msec) {
    struct itimerval timer;

    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;
    timer.it_value.tv_usec = 1000 * (msec % 1000);
    timer.it_interval.tv_sec = 0;
    timer.it_interval.tv_usec = 0;

    setitimer(ITIMER_REAL, &timer, 0);
}
```

```
static void flush_icache(void* begin, size_t size) {
    uintptr_t p = (uintptr_t)begin & ~15UL;
    for (size_t i = 0; i < size; i += 16) {
        asm("icbi 0,%0" : : "r"(p));
        p += 16;
    }
    asm("isync");
}
```

```
DataHeapType* getDataHeap() {
    static char buf[sizeof(DataHeapType)];
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;
    return _theDataHeap;
}
```

```
DataHeapType* getDataHeap() {
    static char buf[sizeof(DataHeapType)];
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;
    return _theDataHeap;
}
```

```
int main(int argc, char **argv) {
    topFrame = (void**)__builtin_frame_address(0);
    setHandler(Trap::TrapSignal, onTrap);
    setHandler(SIGALRM, onTimer);
    setHandler(SIGSEGV, onFault);
    for(Function* f: functions) {
        f->setTrap();
    }
    setTimer(interval);
    int r = stabilizer_main(argc, argv);
    return r;
}
```

```
static void flush_icache(void* begin, size_t size) {
    uintptr_t p = (uintptr_t)begin & ~15UL;
    for (size_t i = 0; i < size; i += 16) {
        asm("icbi 0,%0" : : "r"(p));
        p += 16;
    }
}
```

```
void setTimer(int msec) {
    struct itimerval timer;

    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;
    timer.it_value.tv_usec = 1000 * (msec % 1000);
    timer.it_interval.tv_sec = 0;
    timer.it_interval.tv_usec = 0;

    setitimer(ITIMER_REAL, &timer, 0);
}
```

Why is **A'** faster than **A** ?

A

A'

```
int main(int argc, char **argv) {
    topFrame = (void**)__builtin_frame_address(0);
    setHandler(Trap::TrapSignal, onTrap);
    setHandler(SIGALRM, onTimer);
    setHandler(SIGSEGV, onFault);
    for(Function* f: functions) {
        f->setTrap();
    }
    setTimer(interval);
    int r = stabilizer_main(argc, argv);
    return r;
}

void setTimer(int msec) {
    struct itimerval timer;

    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;
    timer.it_value.tv_usec = 1000 * (msec % 1000);
    timer.it_interval.tv_sec = 0;
    timer.it_interval.tv_usec = 0;

    setitimer(ITIMER_REAL, &timer, 0);
}
```

```
static void flush_icache(void* begin, size_t size) {
    uintptr_t p = (uintptr_t)begin & ~15UL;
    for (size_t i = 0; i < size; i += 16) {
        asm("icbi 0,%0" : : "r"(p));
        p += 16;
    }
    asm("isync");
}
```

```
DataHeapType* getDataHeap() {
    static char buf[sizeof(DataHeapType)];
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;
    return _theDataHeap;
}
```

map to
same set

```
DataHeapType* getDataHeap() {
    static char buf[sizeof(DataHeapType)];
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;
    return _theDataHeap;
}
```

```
int main(int argc, char **argv) {
    topFrame = (void**)__builtin_frame_address(0);
    setHandler(Trap::TrapSignal, onTrap);
    setHandler(SIGALRM, onTimer);
    setHandler(SIGSEGV, onFault);
    for(Function* f: functions) {
        f->setTrap();
    }
    setTimer(interval);
    int r = stabilizer_main(argc, argv);
    return r;
}
```

```
static void flush_icache(void* begin, size_t size) {
    uintptr_t p = (uintptr_t)begin & ~15UL;
    for (size_t i = 0; i < size; i += 16) {
        asm("icbi 0,%0" : : "r"(p));
        p += 16;
    }
}
```

```
void setTimer(int msec) {
    struct itimerval timer;

    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;
    timer.it_value.tv_usec = 1000 * (msec % 1000);
    timer.it_interval.tv_sec = 0;
    timer.it_interval.tv_usec = 0;

    setitimer(ITIMER_REAL, &timer, 0);
}
```

Why is **A'** faster than **A** ?

A

A'

```
int main(int argc, char **argv) {
    topFrame = (void**)__builtin_frame_address(0);
    setHandler(Trap::TrapSignal, onTrap);
    setHandler(SIGALRM, onTimer);
    setHandler(SIGSEGV, onFault);
    for(Function* f: functions) {
        f->setTrap();
    }
    setTimer(interval);
    int r = stabilizer_main(argc, argv);
    return r;
}

void setTimer(int msec) {
    struct itimerval timer;

    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;
    timer.it_value.tv_usec = 1000 * (msec % 1000);
    timer.it_interval.tv_sec = 0;
    timer.it_interval.tv_usec = 0;

    setitimer(ITIMER_REAL, &timer, 0);
}
```

```
static void flush_icache(void* begin, size_t size) {
    uintptr_t p = (uintptr_t)begin & ~15UL;
    for (size_t i = 0; i < size; i += 16) {
        asm("icbi 0,%0" : : "r"(p));
        p += 16;
    }
    asm("isync");
}
```

```
DataHeapType* getDataHeap() {
    static char buf[sizeof(DataHeapType)];
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;
    return _theDataHeap;
}
```

map to
same set

Nothing here

```
DataHeapType* getDataHeap() {
    static char buf[sizeof(DataHeapType)];
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;
    return _theDataHeap;
}
```

```
int main(int argc, char **argv) {
    topFrame = (void**)__builtin_frame_address(0);
    setHandler(Trap::TrapSignal, onTrap);
    setHandler(SIGALRM, onTimer);
    setHandler(SIGSEGV, onFault);
    for(Function* f: functions) {
        f->setTrap();
    }
    setTimer(interval);
    int r = stabilizer_main(argc, argv);
    return r;
}
```




```
static void flush_icache(void* begin, size_t size) {
    uintptr_t p = (uintptr_t)begin & ~15UL;
    for (size_t i = 0; i < size; i += 16) {
        asm("icbi 0,%0" : : "r"(p));
        p += 16;
    }
}
```

```
void setTimer(int msec) {
    struct itimerval timer;

    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;
    timer.it_value.tv_usec = 1000 * (msec % 1000);
    timer.it_interval.tv_sec = 0;
    timer.it_interval.tv_usec = 0;

    setitimer(ITIMER_REAL, &timer, 0);
}
```

Why is **A'** faster than **A** ?



```
int main(int argc, char **argv) {
    topFrame = (void**)__builtin_frame_address(0);
    setHandler(Trap::TrapSignal, onTrap);
    setHandler(SIGALRM, onTimer);
    setHandler(SIGSEGV, onFault);
    for(Function* f: functions) {
        f->setTrap();
    }
    setTimer(interval);
    int r = stabilizer_main(argc, argv);
    return r;
}

void setTimer(int msec) {
    struct itimerval timer;

    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;
    timer.it_value.tv_usec = 1000 * (msec % 1000);
    timer.it_interval.tv_sec = 0;
    timer.it_interval.tv_usec = 0;

    setitimer(ITIMER_REAL, &timer, 0);
}
```

```
static void flush_icache(void* begin, size_t size) {
    uintptr_t p = (uintptr_t)begin & ~15UL;
    for (size_t i = 0; i < size; i += 16) {
        asm("icbi 0,%0" : : "r"(p));
        p += 16;
    }
    asm("isync");
}
```

```
DataHeapType* getDataHeap() {
    static char buf[sizeof(DataHeapType)];
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;
    return _theDataHeap;
}
```

map to
same set

Nothing here →
no conflict

```
DataHeapType* getDataHeap() {
    static char buf[sizeof(DataHeapType)];
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;
    return _theDataHeap;
}
```

```
int main(int argc, char **argv) {
    topFrame = (void**)__builtin_frame_address(0);
    setHandler(Trap::TrapSignal, onTrap);
    setHandler(SIGALRM, onTimer);
    setHandler(SIGSEGV, onFault);
    for(Function* f: functions) {
        f->setTrap();
    }
    setTimer(interval);
    int r = stabilizer_main(argc, argv);
    return r;
}
```



```
static void flush_icache(void* begin, size_t size) {
    uintptr_t p = (uintptr_t)begin & ~15UL;
    for (size_t i = 0; i < size; i += 16) {
        asm("icbi 0,%0" : : "r"(p));
        p += 16;
    }
}
```

```
void setTimer(int msec) {
    struct itimerval timer;

    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;
    timer.it_value.tv_usec = 1000 * (msec % 1000);
    timer.it_interval.tv_sec = 0;
    timer.it_interval.tv_usec = 0;

    setitimer(ITIMER_REAL, &timer, 0);
}
```

It's the cache



```
int main(int argc, char **argv) {
    topFrame = (void**)__builtin_frame_address(0);
    setHandler(Trap::TrapSignal, onTrap);
    setHandler(SIGALRM, onTimer);
    setHandler(SIGSEGV, onFault);
    for(Function* f: functions) {
        f->setTrap();
    }
    setTimer(interval);
    int r = stabilizer_main(argc, argv);
    return r;
}

void setTimer(int msec) {
    struct itimerval timer;

    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;
    timer.it_value.tv_usec = 1000 * (msec % 1000);
    timer.it_interval.tv_sec = 0;
    timer.it_interval.tv_usec = 0;

    setitimer(ITIMER_REAL, &timer, 0);
}
```

```
static void flush_icache(void* begin, size_t size) {
    uintptr_t p = (uintptr_t)begin & ~15UL;
    for (size_t i = 0; i < size; i += 16) {
        asm("icbi 0,%0" : : "r"(p));
        p += 16;
    }
    asm("isync");
}
```

```
DataHeapType* getDataHeap() {
    static char buf[sizeof(DataHeapType)];
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;
    return _theDataHeap;
}
```

```
DataHeapType* getDataHeap() {
    static char buf[sizeof(DataHeapType)];
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;
    return _theDataHeap;
}
```

```
int main(int argc, char **argv) {
    topFrame = (void**)__builtin_frame_address(0);
    setHandler(Trap::TrapSignal, onTrap);
    setHandler(SIGALRM, onTimer);
    setHandler(SIGSEGV, onFault);
    for(Function* f: functions) {
        f->setTrap();
    }
    setTimer(interval);
    int r = stabilizer_main(argc, argv);
    return r;
}
```

```
static void flush_icache(void* begin, size_t size) {
    uintptr_t p = (uintptr_t)begin & ~15UL;
    for (size_t i = 0; i < size; i += 16) {
        asm("icbi 0,%0" : : "r"(p));
        p += 16;
    }
}
```

```
void setTimer(int msec) {
    struct itimerval timer;

    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;
    timer.it_value.tv_usec = 1000 * (msec % 1000);
    timer.it_interval.tv_sec = 0;
    timer.it_interval.tv_usec = 0;

    setitimer(ITIMER_REAL, &timer, 0);
}
```

It's the cache

or branch predictor



```
int main(int argc, char **argv) {
    topFrame = (void**)__builtin_frame_address(0);
    setHandler(Trap::TrapSignal, onTrap);
    setHandler(SIGALRM, onTimer);
    setHandler(SIGSEGV, onFault);
    for(Function* f: functions) {
        f->setTrap();
    }
    setTimer(interval);
    int r = stabilizer_main(argc, argv);
    return r;
}

void setTimer(int msec) {
    struct itimerval timer;

    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;
    timer.it_value.tv_usec = 1000 * (msec % 1000);
    timer.it_interval.tv_sec = 0;
    timer.it_interval.tv_usec = 0;

    setitimer(ITIMER_REAL, &timer, 0);
}
```

```
static void flush_icache(void* begin, size_t size) {
    uintptr_t p = (uintptr_t)begin & ~15UL;
    for (size_t i = 0; i < size; i += 16) {
        asm("icbi 0,%0" : : "r"(p));
        p += 16;
    }
    asm("isync");
}
```

```
DataHeapType* getDataHeap() {
    static char buf[sizeof(DataHeapType)];
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;
    return _theDataHeap;
}
```

```
DataHeapType* getDataHeap() {
    static char buf[sizeof(DataHeapType)];
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;
    return _theDataHeap;
}
```

```
int main(int argc, char **argv) {
    topFrame = (void**)__builtin_frame_address(0);
    setHandler(Trap::TrapSignal, onTrap);
    setHandler(SIGALRM, onTimer);
    setHandler(SIGSEGV, onFault);
    for(Function* f: functions) {
        f->setTrap();
    }
    setTimer(interval);
    int r = stabilizer_main(argc, argv);
    return r;
}
```

```
static void flush_icache(void* begin, size_t size) {
    uintptr_t p = (uintptr_t)begin & ~15UL;
    for (size_t i = 0; i < size; i += 16) {
        asm("icbi 0,%0" : : "r"(p));
        p += 16;
    }
}
```

```
void setTimer(int msec) {
    struct itimerval timer;

    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;
    timer.it_value.tv_usec = 1000 * (msec % 1000);
    timer.it_interval.tv_sec = 0;
    timer.it_interval.tv_usec = 0;


    setitimer(ITIMER_REAL, &timer, 0);
}
```

It's the cache

or TLB



or branch predictor



```
int main(int argc, char **argv) {
    topFrame = (void**)__builtin_frame_address(0);
    setHandler(Trap::TrapSignal, onTrap);
    setHandler(SIGALRM, onTimer);
    setHandler(SIGSEGV, onFault);
    for(Function* f: functions) {
        f->setTrap();
    }
    setTimer(interval);
    int r = stabilizer_main(argc, argv);
    return r;
}

void setTimer(int msec) {
    struct itimerval timer;

    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;
    timer.it_value.tv_usec = 1000 * (msec % 1000);
    timer.it_interval.tv_sec = 0;
    timer.it_interval.tv_usec = 0;

    setitimer(ITIMER_REAL, &timer, 0);
}
```

```
static void flush_icache(void* begin, size_t size) {
    uintptr_t p = (uintptr_t)begin & ~15UL;
    for (size_t i = 0; i < size; i += 16) {
        asm("icbi 0,%0" : : "r"(p));
        p += 16;
    }
    asm("isync");
}
```

```
DataHeapType* getDataHeap() {
    static char buf[sizeof(DataHeapType)];
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;
    return _theDataHeap;
}
```

```
DataHeapType* getDataHeap() {
    static char buf[sizeof(DataHeapType)];
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;
    return _theDataHeap;
}
```

```
int main(int argc, char **argv) {
    topFrame = (void**)__builtin_frame_address(0);
    setHandler(Trap::TrapSignal, onTrap);
    setHandler(SIGALRM, onTimer);
    setHandler(SIGSEGV, onFault);
    for(Function* f: functions) {
        f->setTrap();
    }
    setTimer(interval);
    int r = stabilizer_main(argc, argv);
    return r;
}
```

```
static void flush_icache(void* begin, size_t size) {
    uintptr_t p = (uintptr_t)begin & ~15UL;
    for (size_t i = 0; i < size; i += 16) {
        asm("icbi 0,%0" : : "r"(p));
        p += 16;
    }
}
```


```
void setTimer(int msec) {
    struct itimerval timer;

    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;
    timer.it_value.tv_usec = 1000 * (msec % 1000);
    timer.it_interval.tv_sec = 0;
    timer.it_interval.tv_usec = 0;

    setitimer(ITIMER_REAL, &timer, 0);
}
```

It's the cache

or TLB



or branch predictor or branch



target predictor

```
int main(int argc, char **argv) {
    topFrame = (void**)__builtin_frame_address(0);
    setHandler(Trap::TrapSignal, onTrap);
    setHandler(SIGALRM, onTimer);
    setHandler(SIGSEGV, onFault);
    for(Function* f: functions) {
        f->setTrap();
    }
    setTimer(interval);
    int r = stabilizer_main(argc, argv);
    return r;
}

void setTimer(int msec) {
    struct itimerval timer;

    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;
    timer.it_value.tv_usec = 1000 * (msec % 1000);
    timer.it_interval.tv_sec = 0;
    timer.it_interval.tv_usec = 0;

    setitimer(ITIMER_REAL, &timer, 0);
}
}
```

```
static void flush_icache(void* begin, size_t size) {
    uintptr_t p = (uintptr_t)begin & ~15UL;
    for (size_t i = 0; i < size; i += 16) {
        asm("icbi 0,%0" : : "r"(p));
        p += 16;
    }
    asm("isync");
}
}
```

```
DataHeapType* getDataHeap() {
    static char buf[sizeof(DataHeapType)];
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;
    return _theDataHeap;
}
}
```

```
DataHeapType* getDataHeap() {
    static char buf[sizeof(DataHeapType)];
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;
    return _theDataHeap;
}
}
```

```
int main(int argc, char **argv) {
    topFrame = (void**)__builtin_frame_address(0);
    setHandler(Trap::TrapSignal, onTrap);
    setHandler(SIGALRM, onTimer);
    setHandler(SIGSEGV, onFault);
    for(Function* f: functions) {
        f->setTrap();
    }
    setTimer(interval);
    int r = stabilizer_main(argc, argv);
    return r;
}
}
```

```
static void flush_icache(void* begin, size_t size) {
    uintptr_t p = (uintptr_t)begin & ~15UL;
    for (size_t i = 0; i < size; i += 16) {
        asm("icbi 0,%0" : : "r"(p));
        p += 16;
    }
}
}
```


```
void setTimer(int msec) {
    struct itimerval timer;

    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;
    timer.it_value.tv_usec = 1000 * (msec % 1000);
    timer.it_interval.tv_sec = 0;
    timer.it_interval.tv_usec = 0;

    setitimer(ITIMER_REAL, &timer, 0);
}
}
```


It's the cache

or TLB



or branch predictor or branch



target predictor or prefetcher

```
int main(int argc, char **argv) {
    topFrame = (void**)__builtin_frame_address(0);
    setHandler(Trap::TrapSignal, onTrap);
    setHandler(SIGALRM, onTimer);
    setHandler(SIGSEGV, onFault);
    for(Function* f: functions) {
        f->setTrap();
    }
    setTimer(interval);
    int r = stabilizer_main(argc, argv);
    return r;
}

void setTimer(int msec) {
    struct itimerval timer;

    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;
    timer.it_value.tv_usec = 1000 * (msec % 1000);
    timer.it_interval.tv_sec = 0;
    timer.it_interval.tv_usec = 0;

    setitimer(ITIMER_REAL, &timer, 0);
}

```

```
static void flush_icache(void* begin, size_t size) {
    uintptr_t p = (uintptr_t)begin & ~15UL;
    for (size_t i = 0; i < size; i += 16) {
        asm("icbi 0,%0" : : "r"(p));
        p += 16;
    }
    asm("isync");
}

```

```
DataHeapType* getDataHeap() {
    static char buf[sizeof(DataHeapType)];
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;
    return _theDataHeap;
}

```

```
DataHeapType* getDataHeap() {
    static char buf[sizeof(DataHeapType)];
    static DataHeapType* _theDataHeap = new (buf) DataHeapType;
    return _theDataHeap;
}

```

```
int main(int argc, char **argv) {
    topFrame = (void**)__builtin_frame_address(0);
    setHandler(Trap::TrapSignal, onTrap);
    setHandler(SIGALRM, onTimer);
    setHandler(SIGSEGV, onFault);
    for(Function* f: functions) {
        f->setTrap();
    }
    setTimer(interval);
    int r = stabilizer_main(argc, argv);
    return r;
}

```

```
static void flush_icache(void* begin, size_t size) {
    uintptr_t p = (uintptr_t)begin & ~15UL;
    for (size_t i = 0; i < size; i += 16) {
        asm("icbi 0,%0" : : "r"(p));
        p += 16;
    }
}

```


```
void setTimer(int msec) {
    struct itimerval timer;

    timer.it_value.tv_sec = (msec - msec % 1000) / 1000;
    timer.it_value.tv_usec = 1000 * (msec % 1000);
    timer.it_interval.tv_sec = 0;
    timer.it_interval.tv_usec = 0;

    setitimer(ITIMER_REAL, &timer, 0);
}


```


Speedup of -0.3 over -0.2

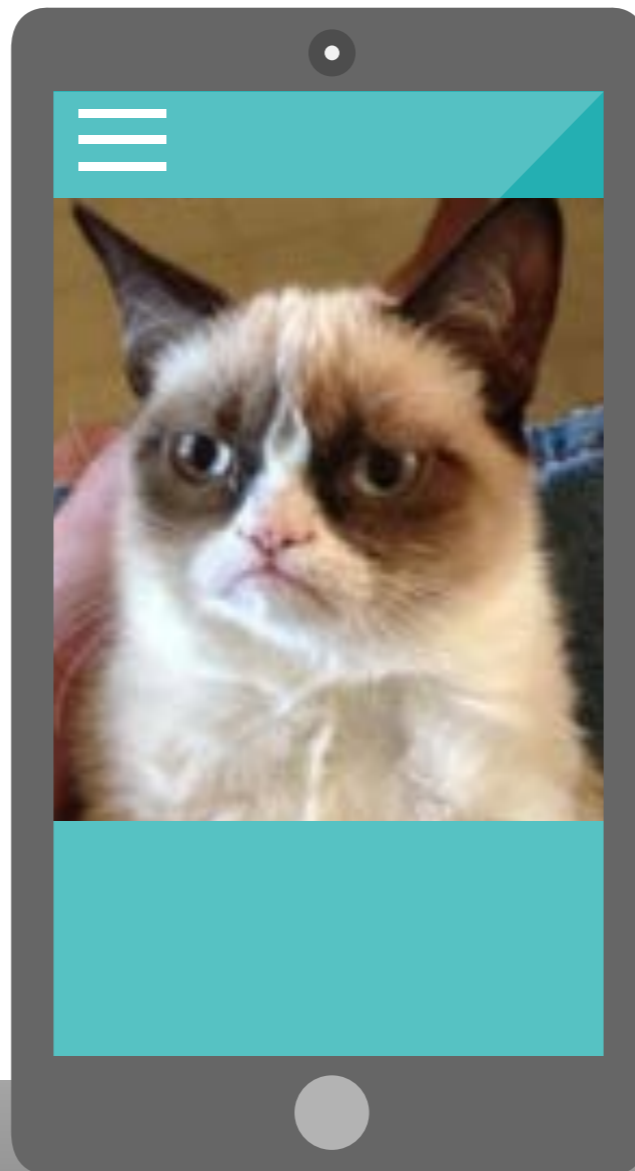




help?

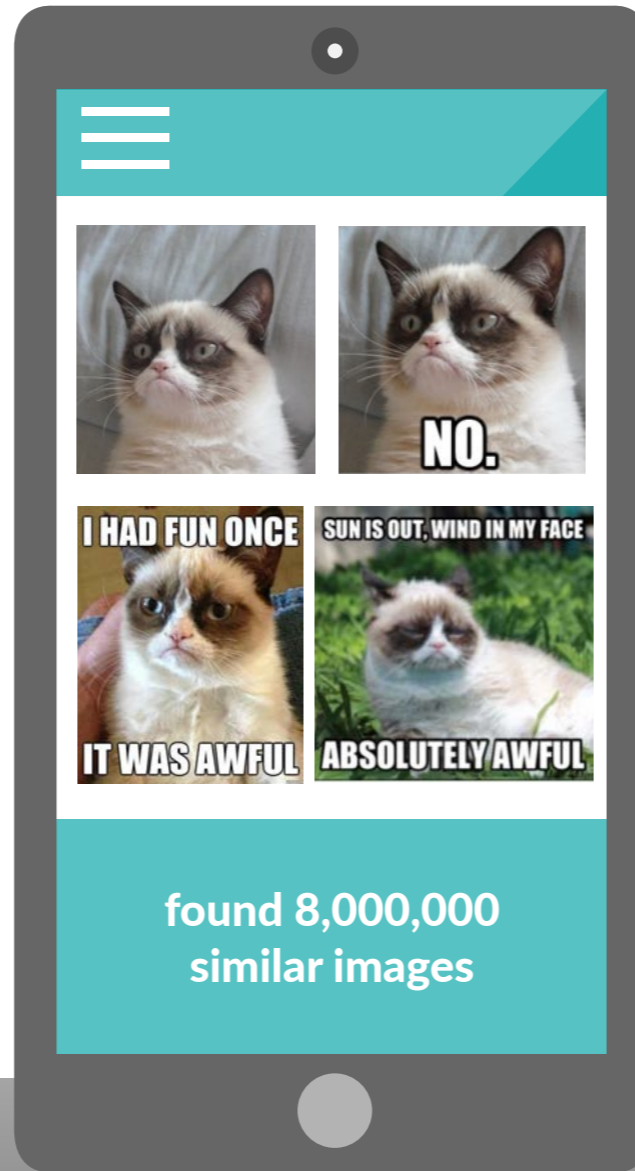


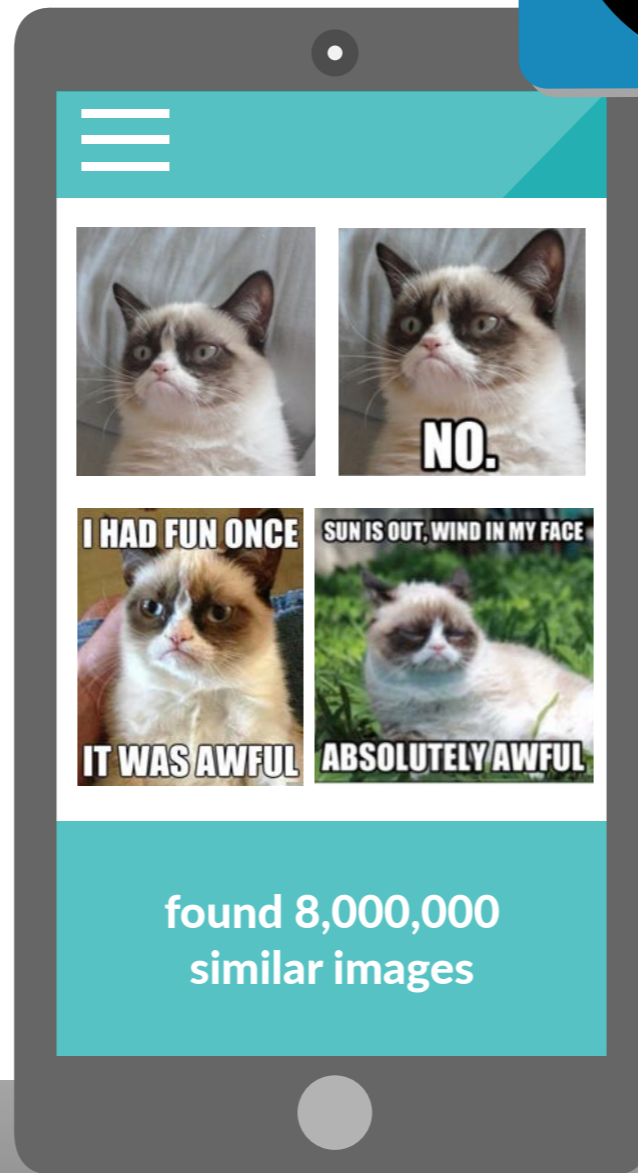




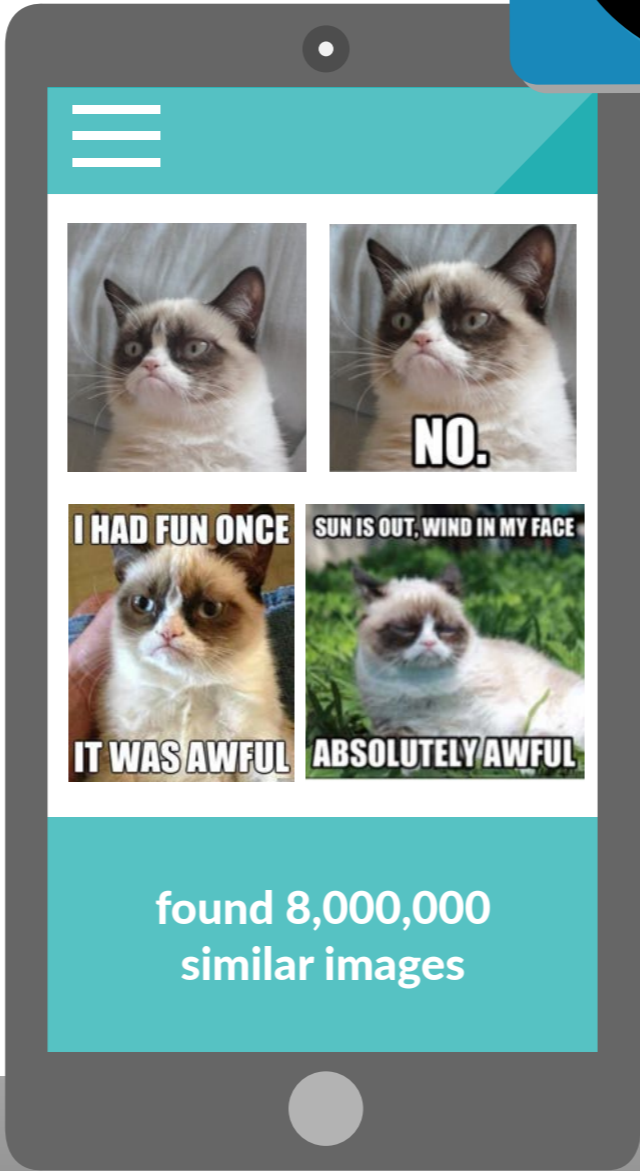









It's going to *totally* disrupt image search.




The Prototype gle




The Prototype gle

Take a picture




The Prototype gle

Take a picture

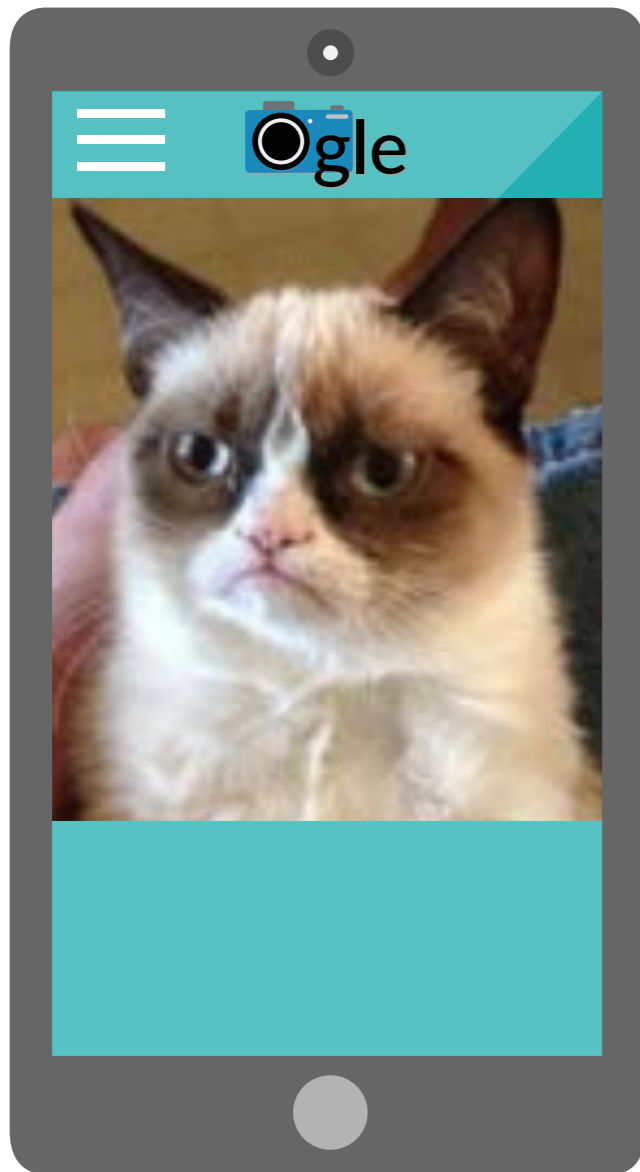


Send it to Ogle





The Prototype gle

Take a picture



Send it to Ogle



Add it to
the database





The Prototype gle

Take a picture




Send it to Ogle



Add it to the database




Find similar pictures



The Prototype gle




Take a picture



Send it to Ogle



Add it to the database




Send results




Find similar pictures







Send results




Find similar pictures




Send results




Find similar pictures







Send results

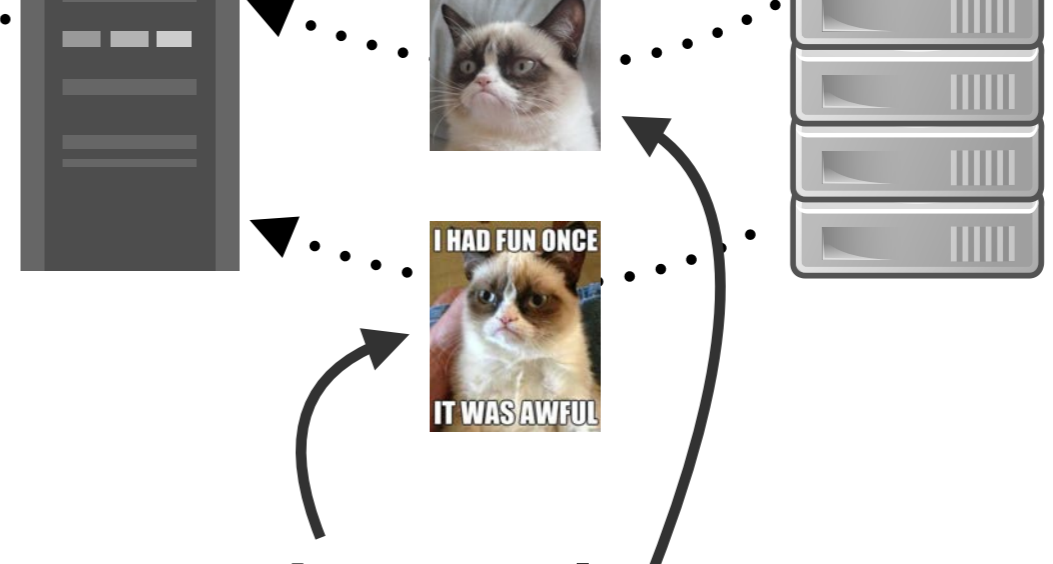


Find similar pictures







Send results

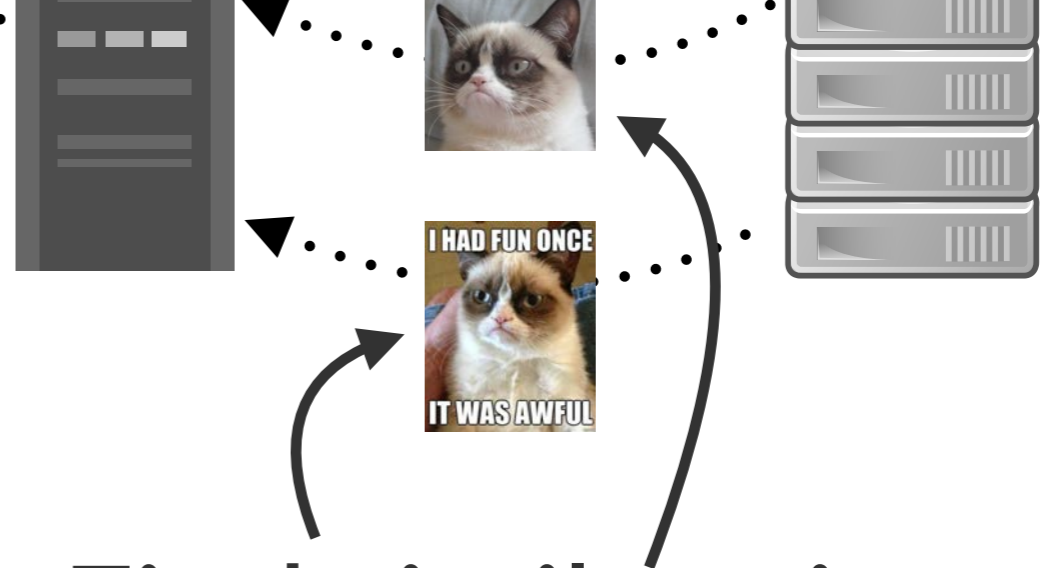


Find similar pictures







Send results

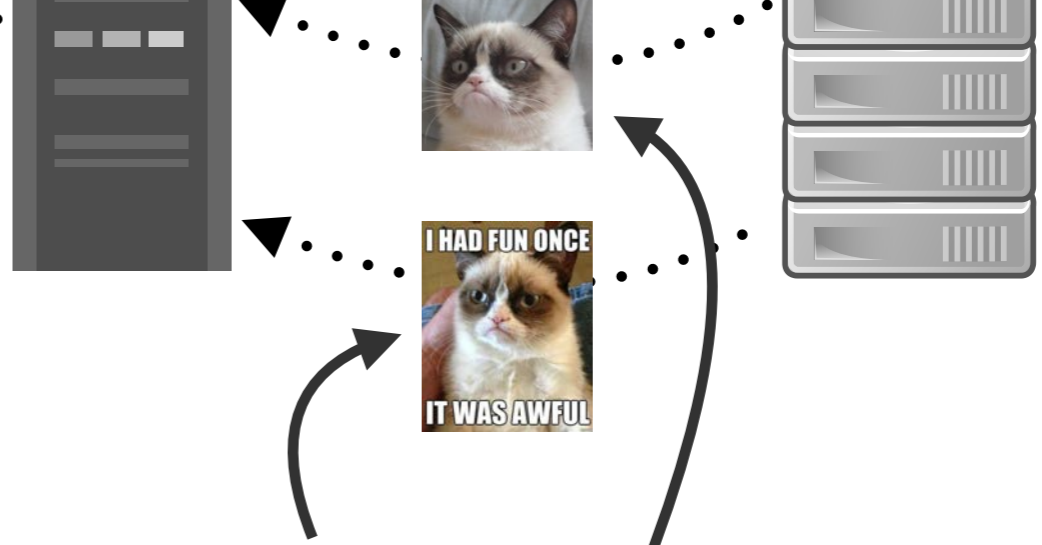


Find similar pictures

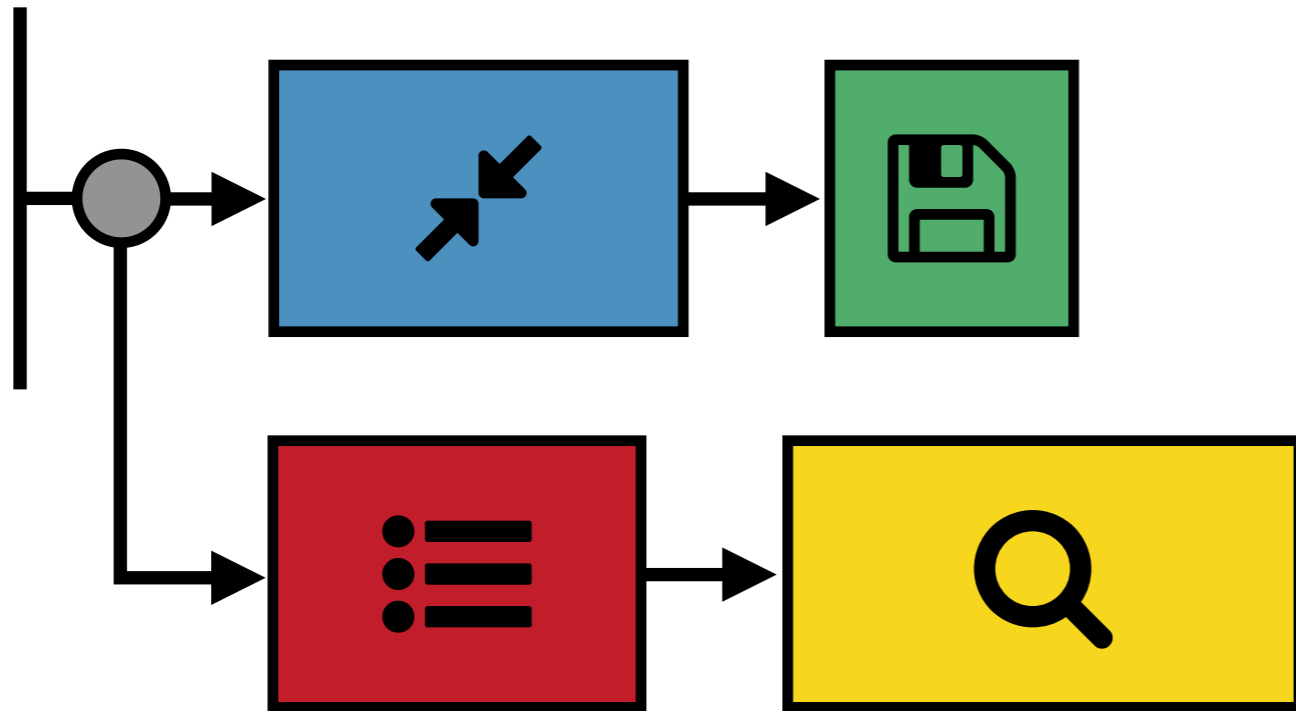





Send results

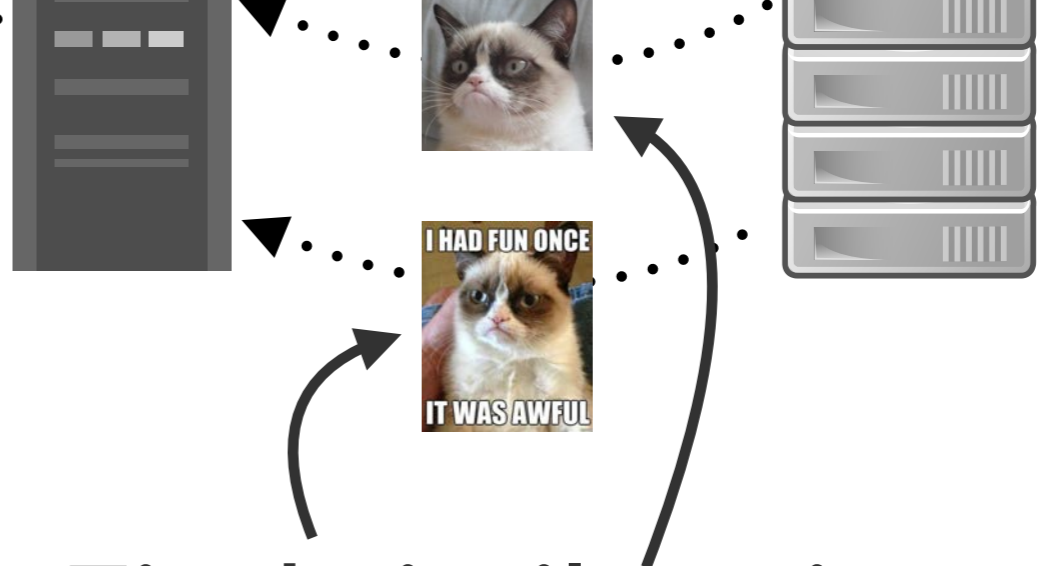


Find similar pictures







Send results




Find similar pictures

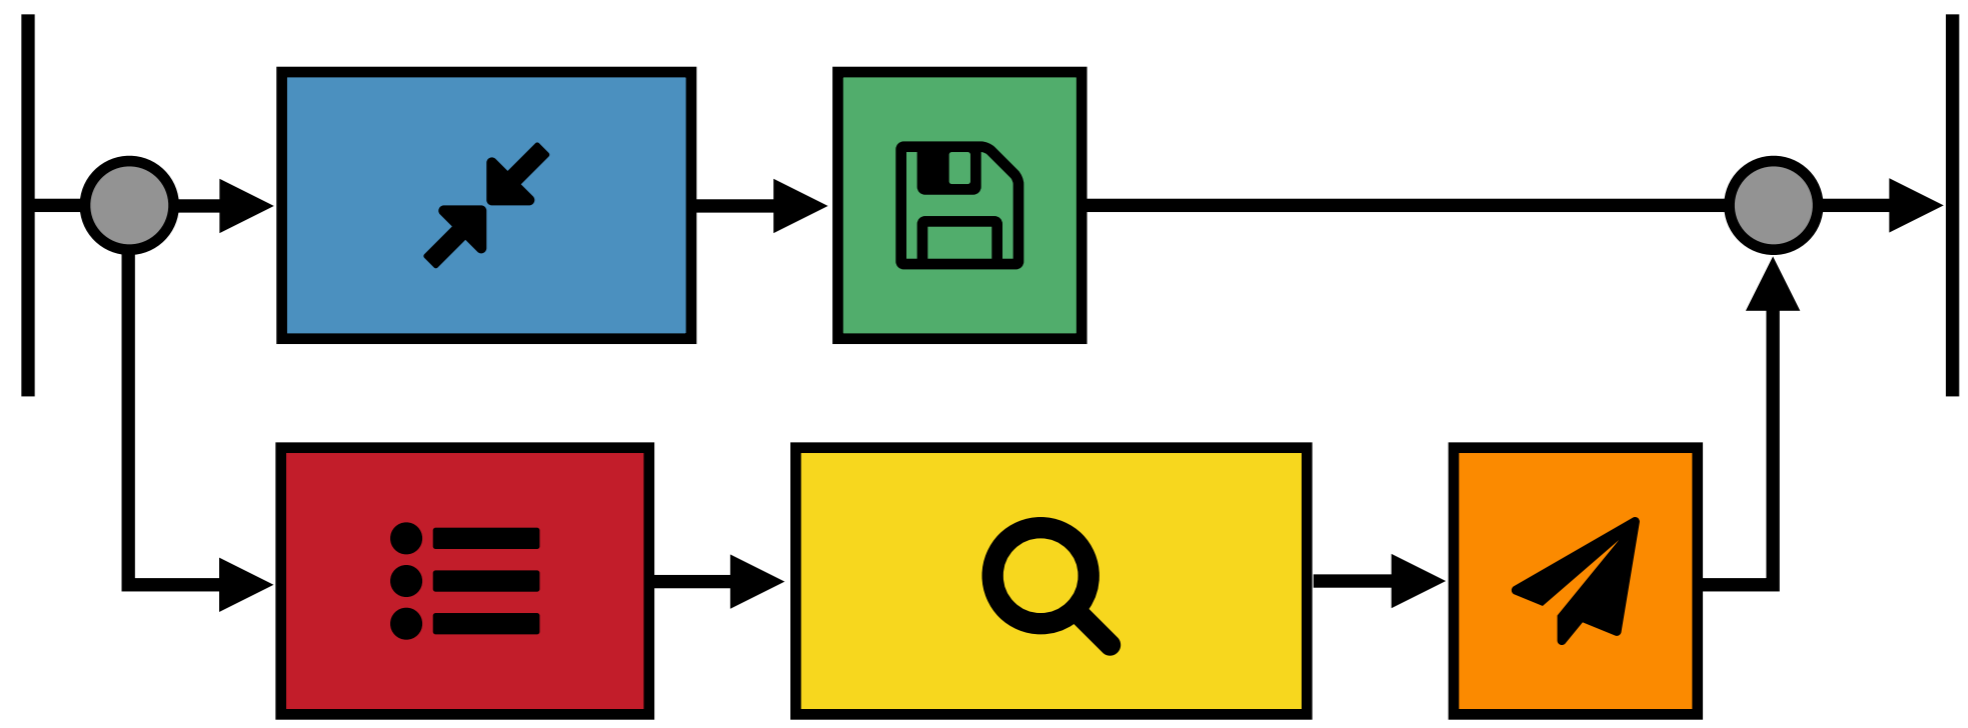





Send results




Find similar pictures

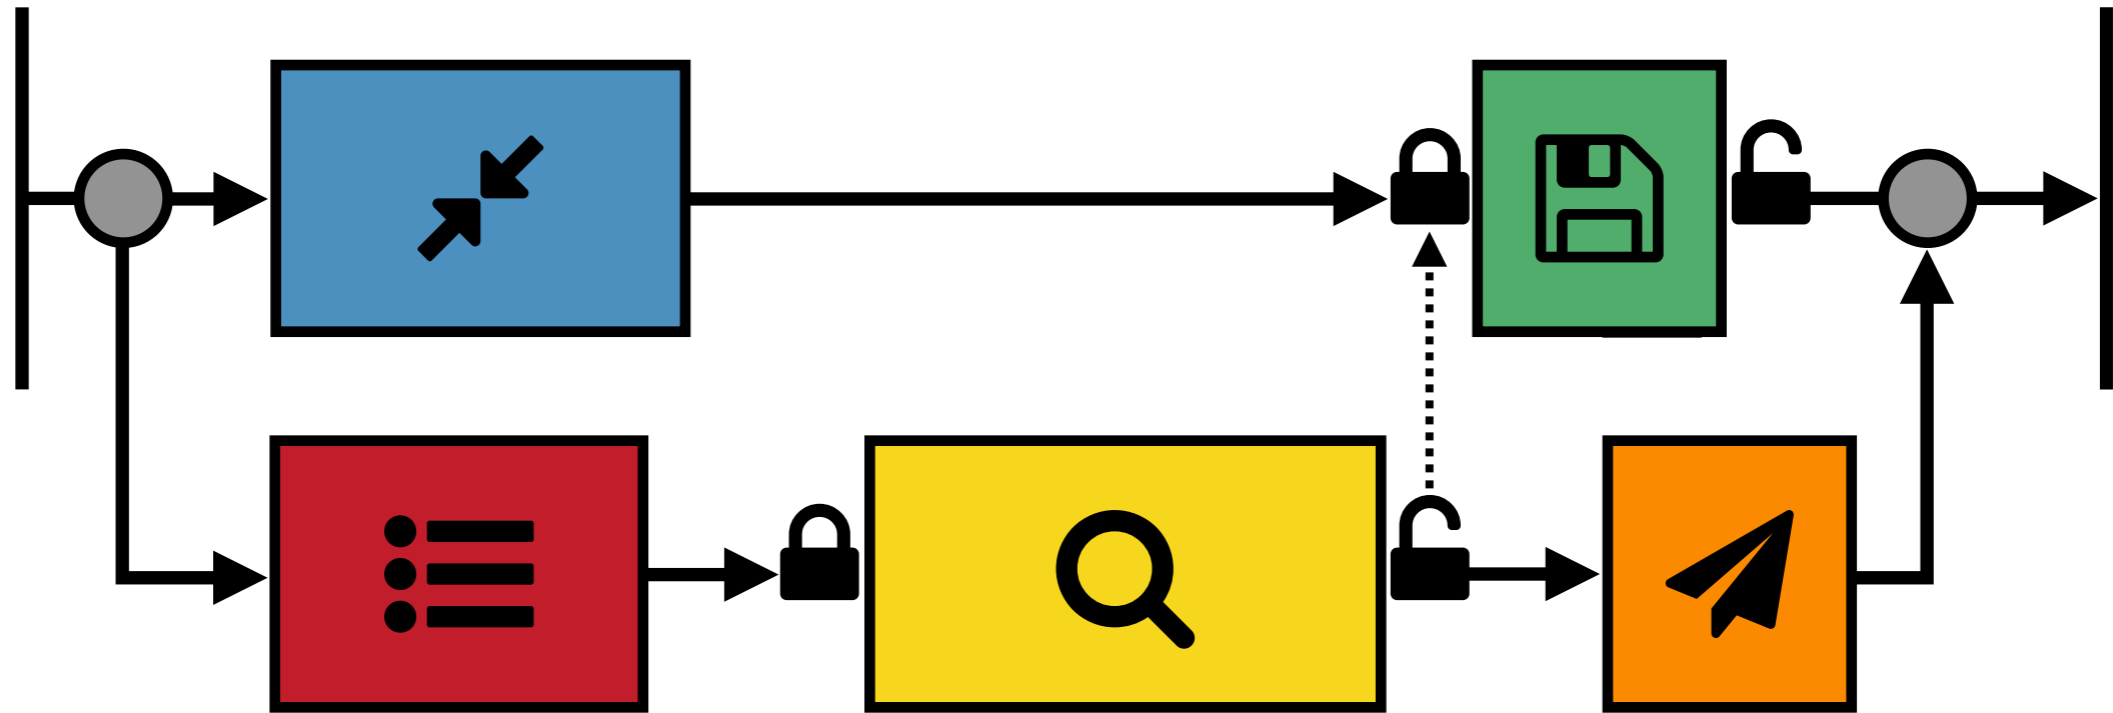


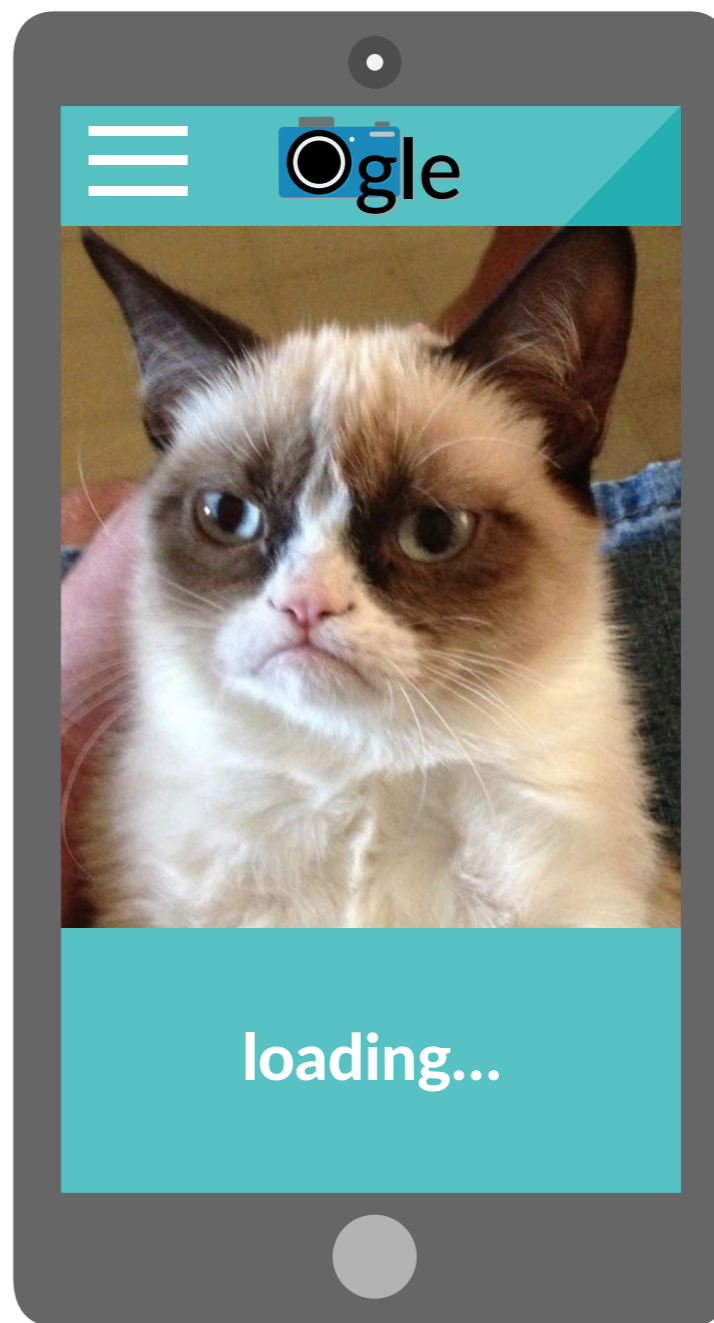


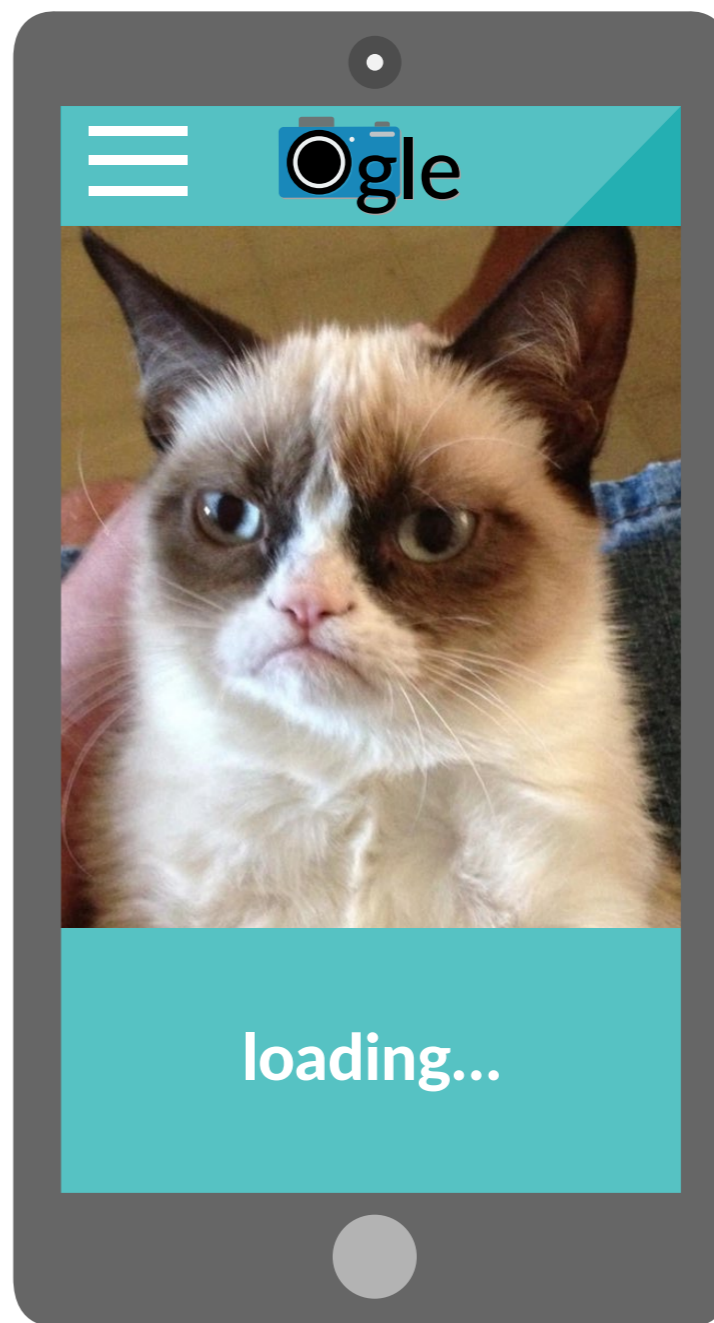
Send results




Find similar pictures









Ogle is too slow!




Software Profilers



Software Profilers




Software Profilers



%	cumulative		
time	seconds	calls	name
20.05	8.02	1	↖↗
9.56	3.82	1	📄
19.95	7.98	1	☰
45.19	11.31	1	🔍
5.25	2.10	1	📧


Software Profilers



**Number of calls
to each function**

%	cumulative	calls	name
time	seconds		
20.05	8.02	1	↖
9.56	3.82	1	📄
19.95	7.98	1	☰
45.19	11.31	1	🔍
5.25	2.10	1	📧

Software Profilers




**Number of calls
to each function**

**Runtime for
each function**

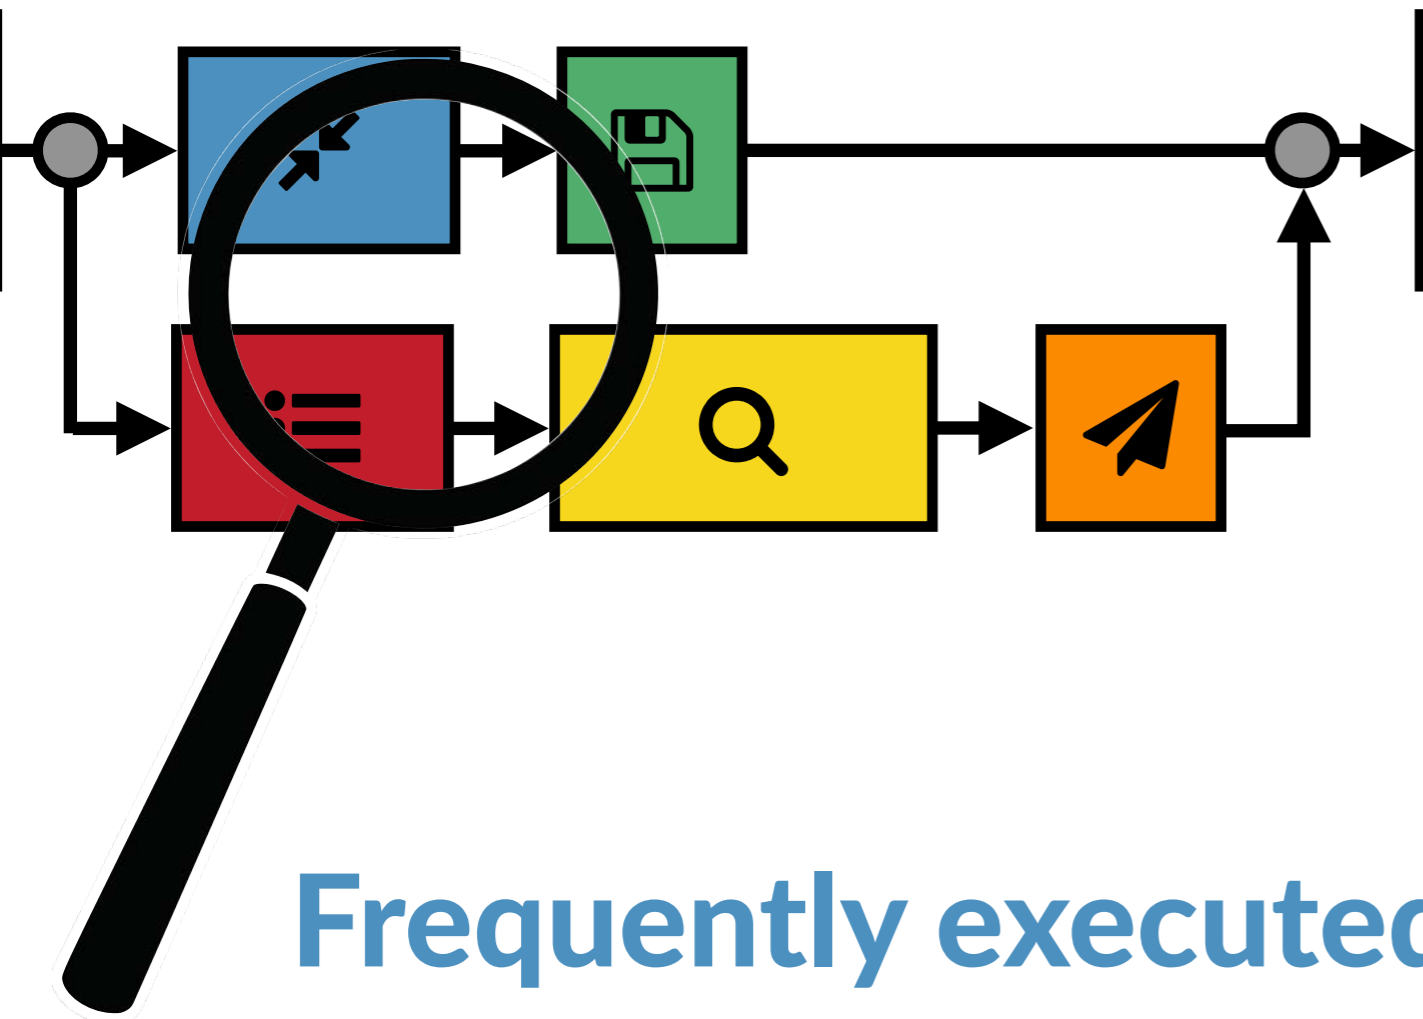
%	cumulative	calls	name
time	seconds		
20.05	8.02	1	↖
9.56	3.82	1	📄
19.95	7.98	1	☰
45.19	11.31	1	🔍
5.25	2.10	1	✈️

Software Profilers



%	cumulative	
time	seconds	ca
20.05	8.02	1
9.56	3.82	1
19.95	7.98	1
45.19	11.31	1
5.25	2.10	1

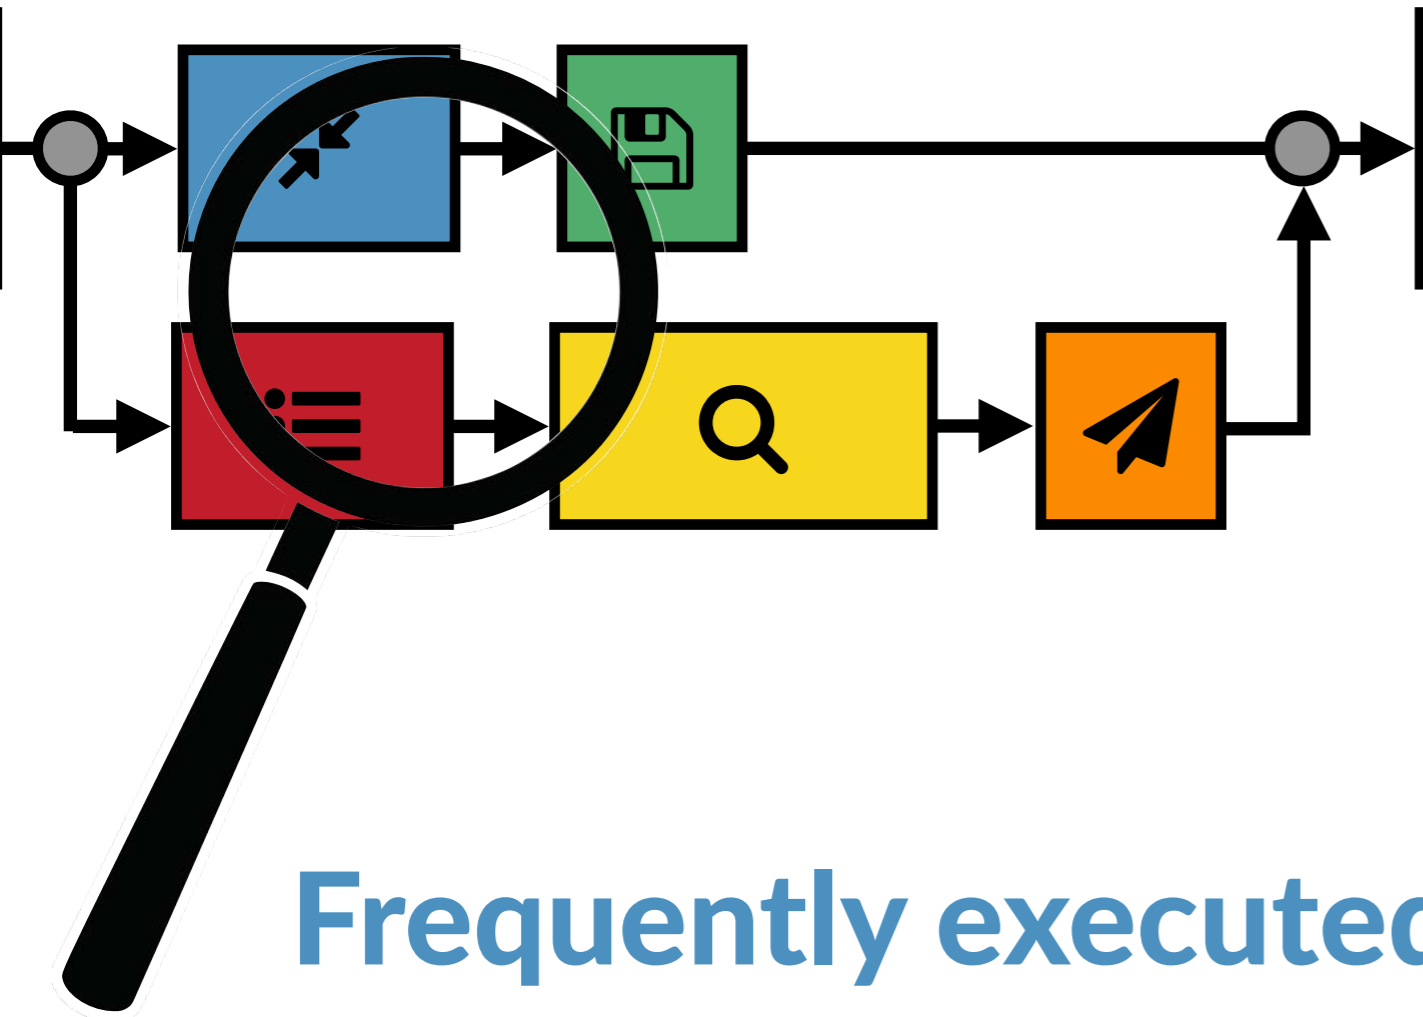
Software Profilers



Frequently executed code

% time	cumulative seconds	ca
20.05	8.02	1
9.56	3.82	1
19.95	7.98	1
45.19	11.31	1
5.25	2.10	1

Software Profilers



Frequently executed code

Code that runs for a long time

% time	cumulative seconds	ca
20.05	8.02	1
9.56	3.82	1
19.95	7.98	1
45.19	11.31	1
5.25	2.10	1

Software Profilers

Are these places where Homer should focus on performance?

Frequently executed code

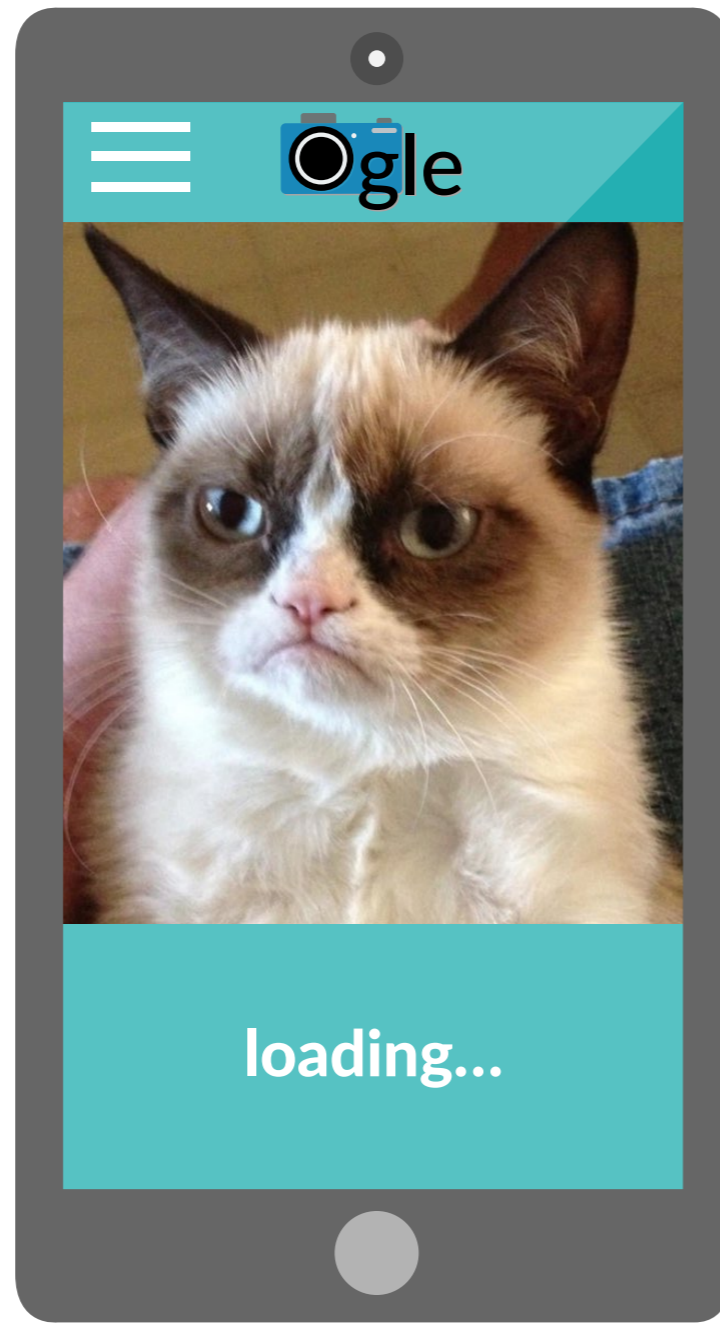
Code that runs for a long time

%	cumulative	
time	seconds	ca
20.05	8.02	1
9.56	3.82	1
19.95	7.98	1
45.19	11.31	1
5.25	2.10	1



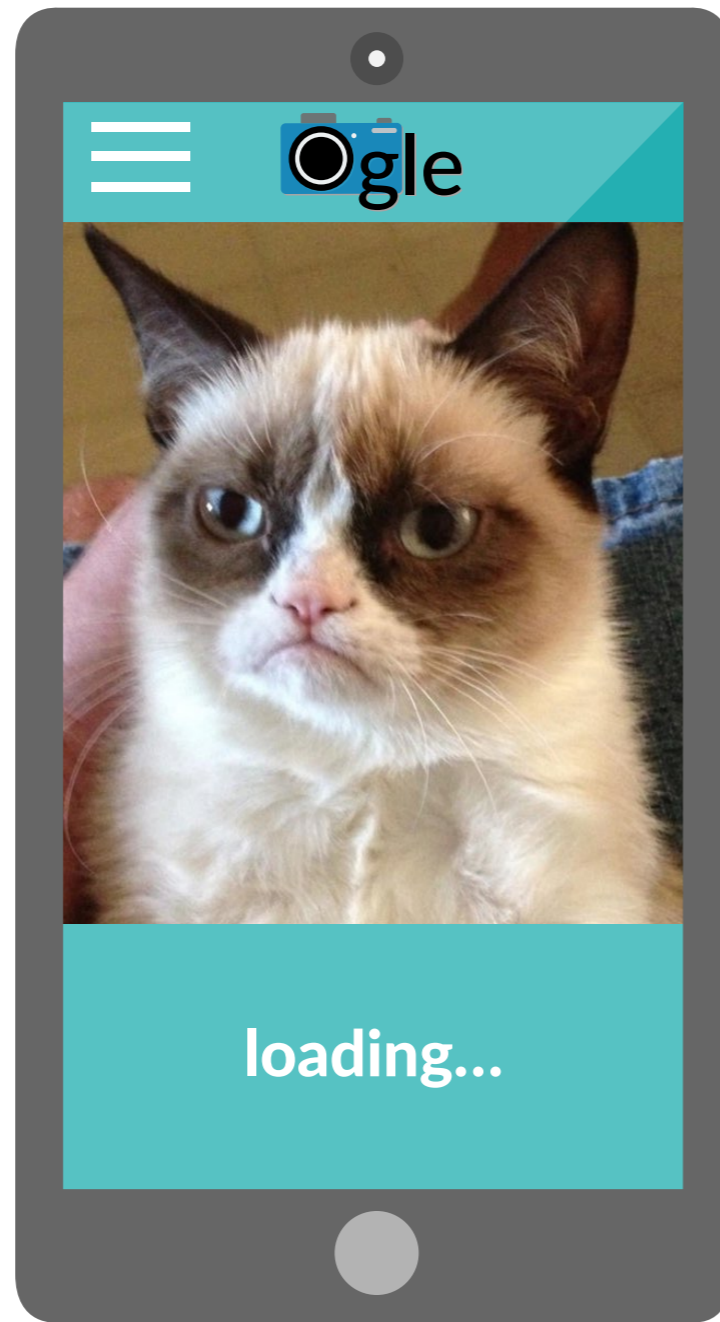
Would this speed up Ogle?





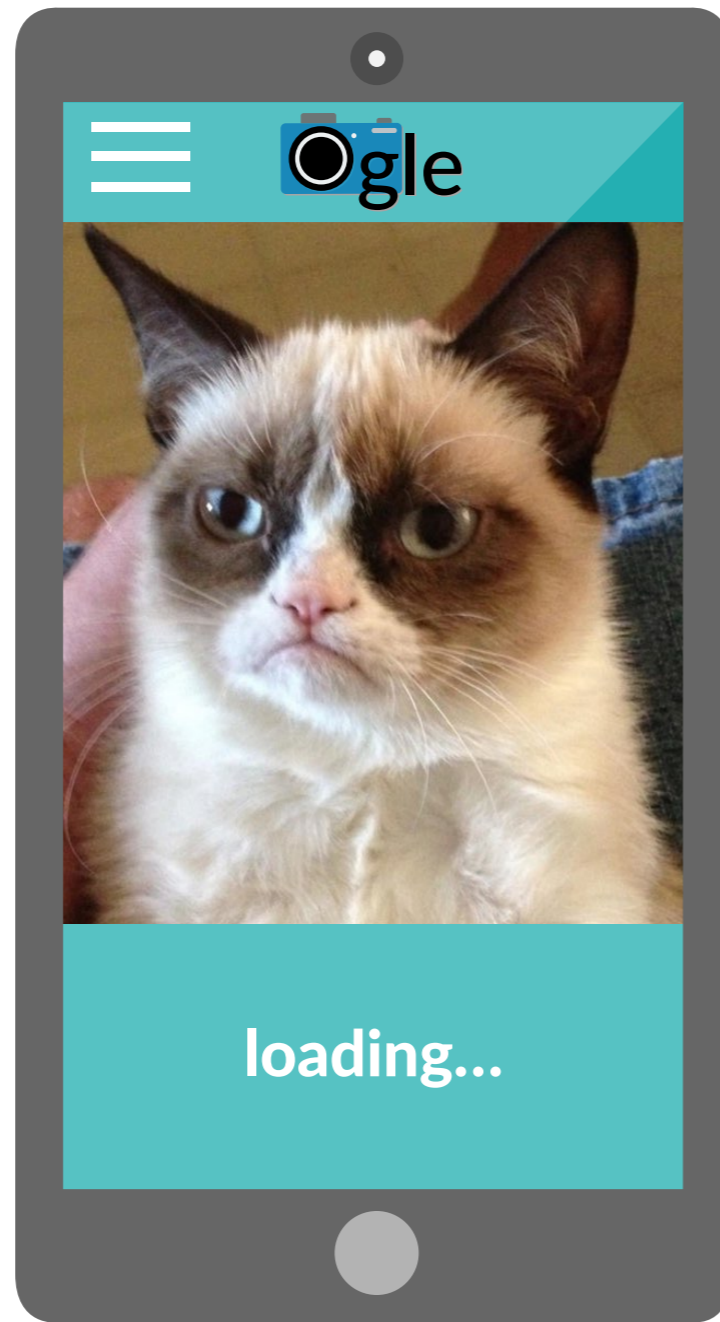
Would this speed up Ogle?




Would this speed up Ogle?



Would this speed up Ogle?




Would this speed up Ogle?



Would this speed up Ogle?


Frequently executed code



Would this speed up Ogle?

Frequently executed code

Code that runs for a long time




Would this speed up Ogle?

Frequently executed code

Code that runs for a long time

Profilers do a bad job finding important code in modern applications!

parallel / asynchronous / concurrent



Would this speed up Ogle?

Frequently executed code


Code that runs for a long time

Profilers do a bad job finding important code in modern applications!


parallel / asynchronous / concurrent




We need to do better!



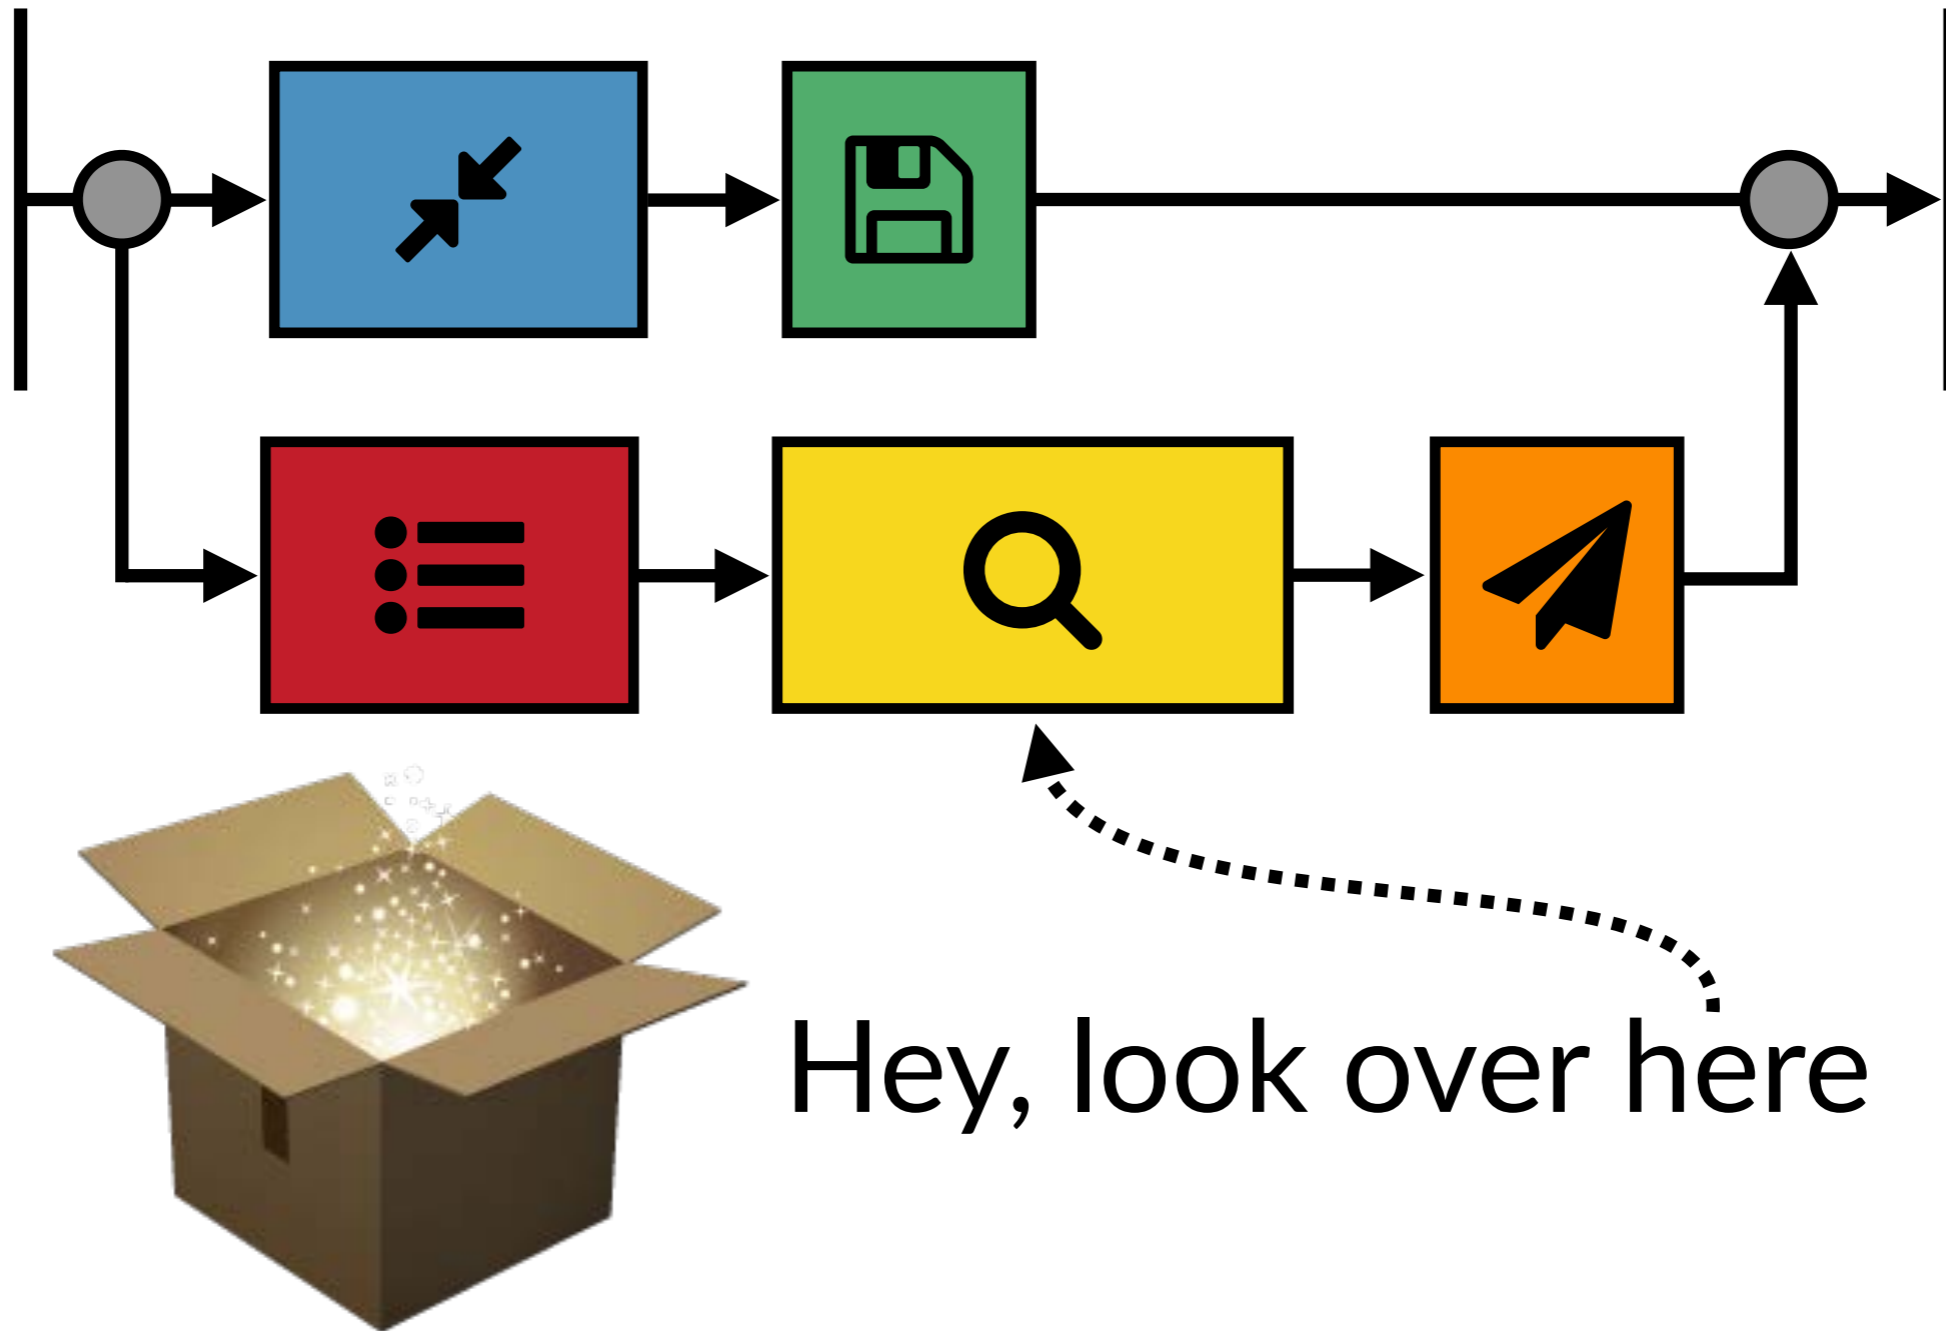
What *would* speed up Ogle?



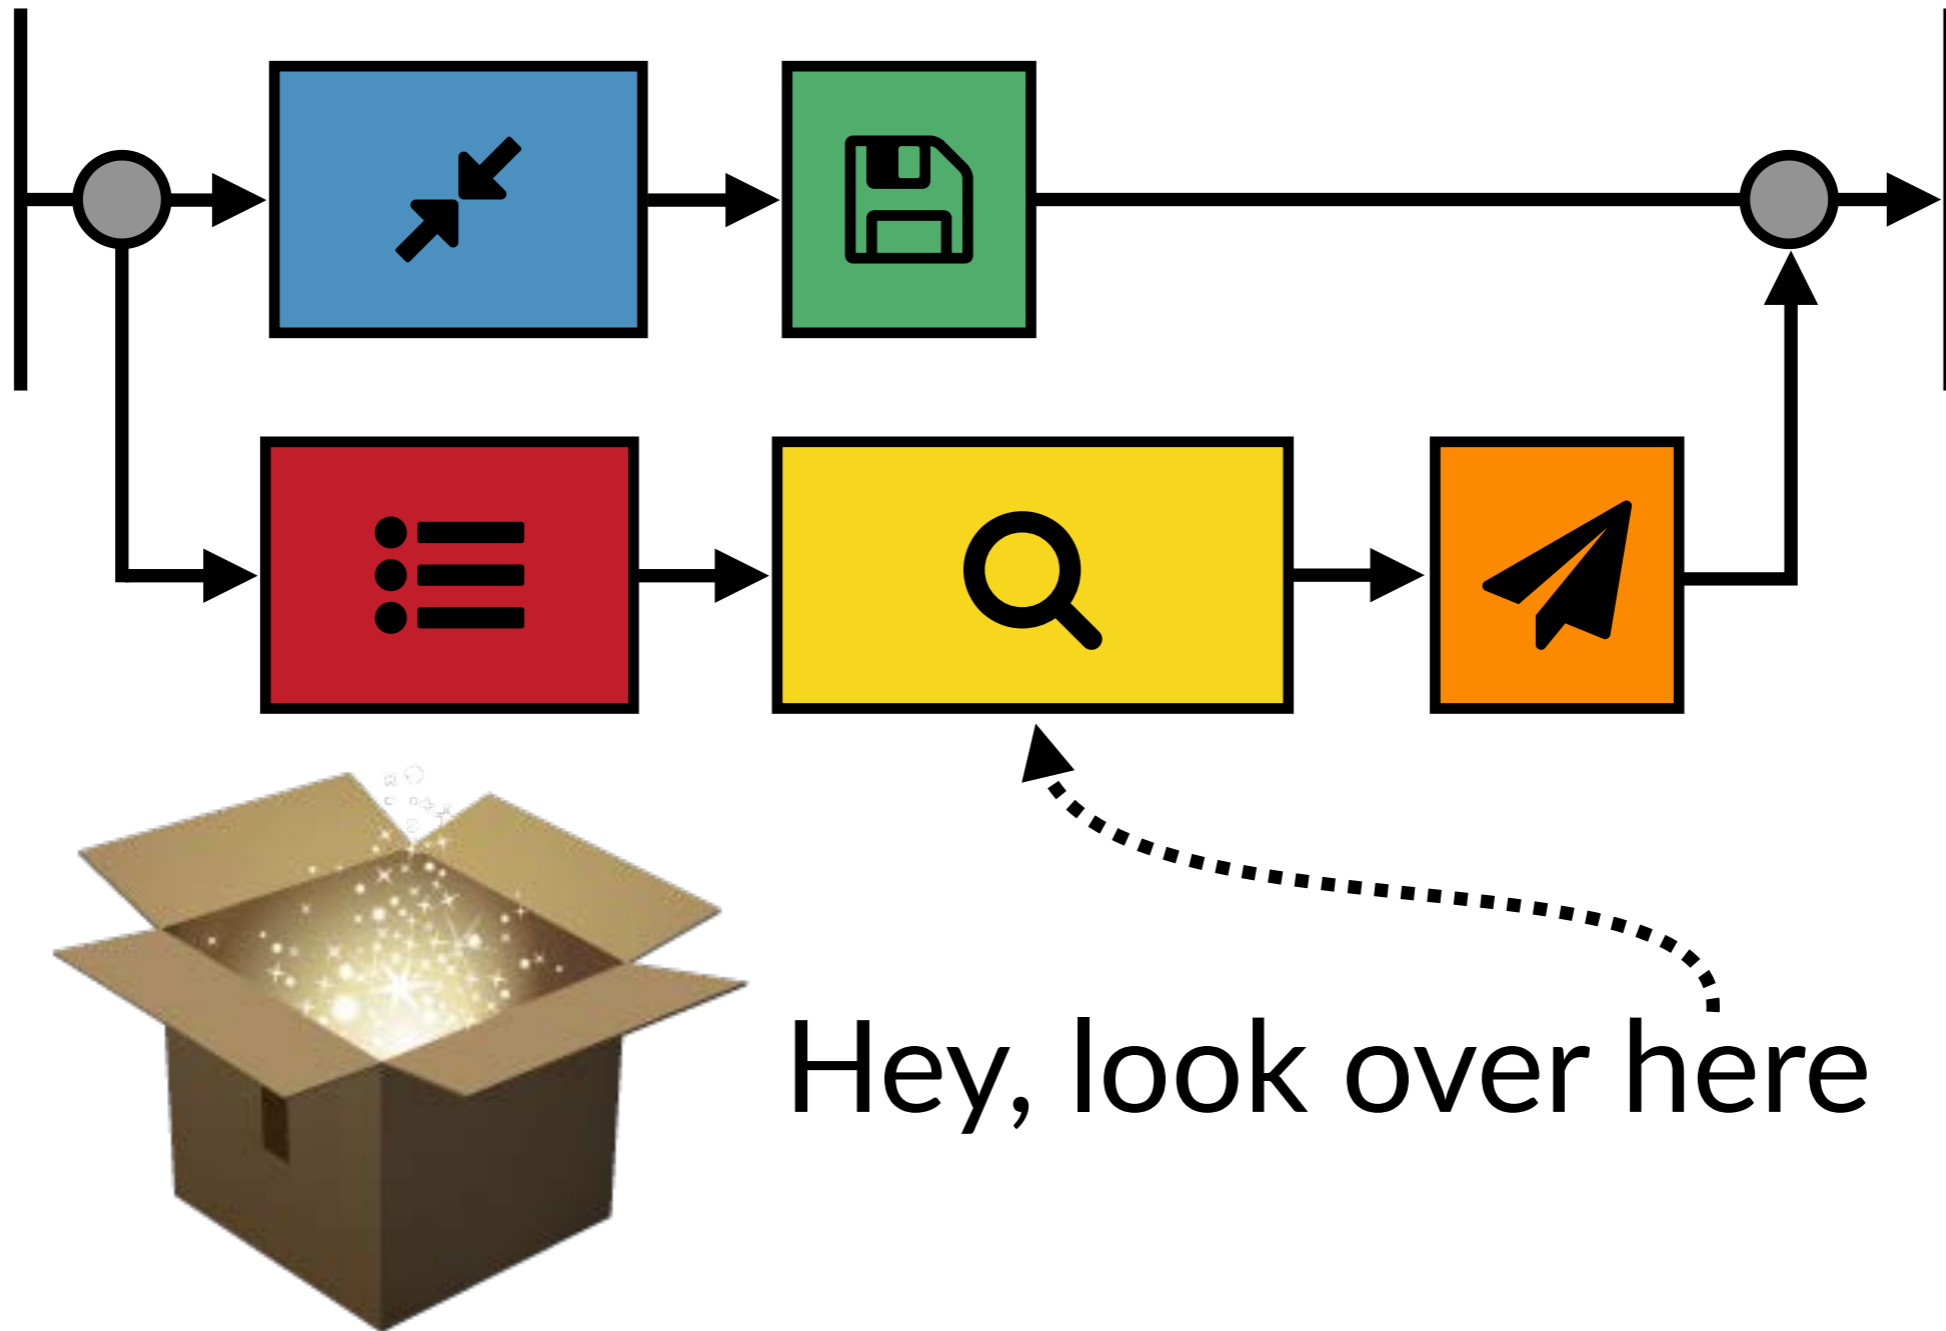
What *would* speed up Ogle?



What *would* speed up Ogle?



What *would* speed up Ogle?



What would this information look like?

Causal Profile

Causal Profile

Tells you where optimizations
will make a difference


Causal Profile

Tells you where optimizations
will make a difference



Causal Profile



Tells you where optimizations will make a difference



Causal Profile

Tells you where optimizations will make a difference


Program Speedup



Speedup


Causal Profile

Tells you where optimizations will make a difference



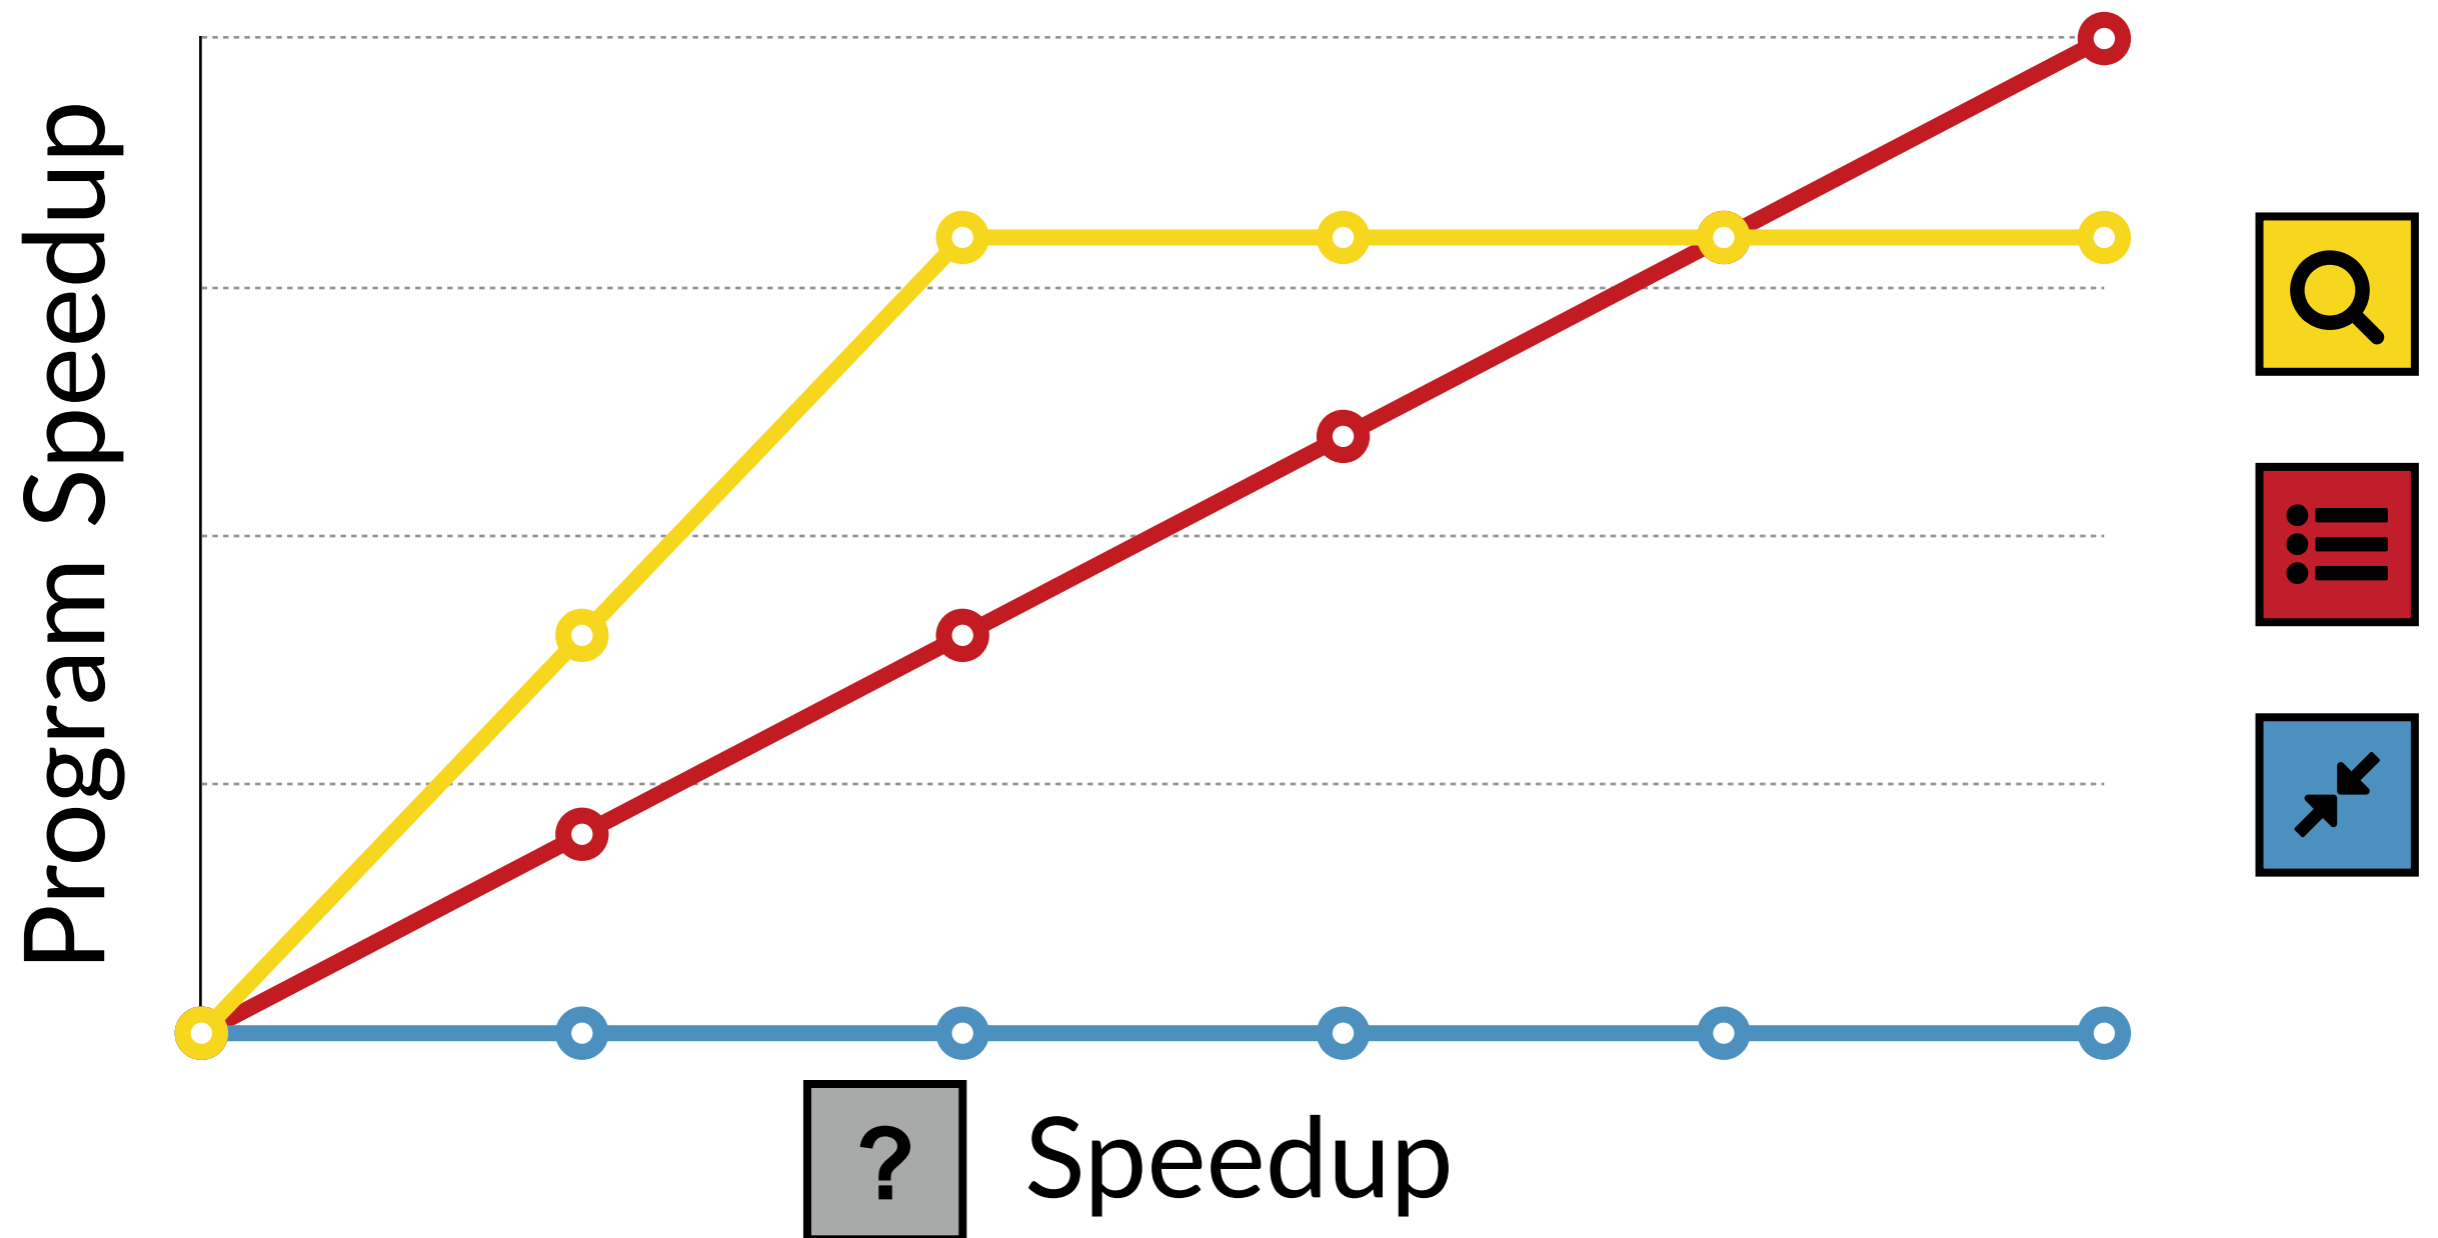
Causal Profile

Tells you where optimizations will make a difference




Causal Profile

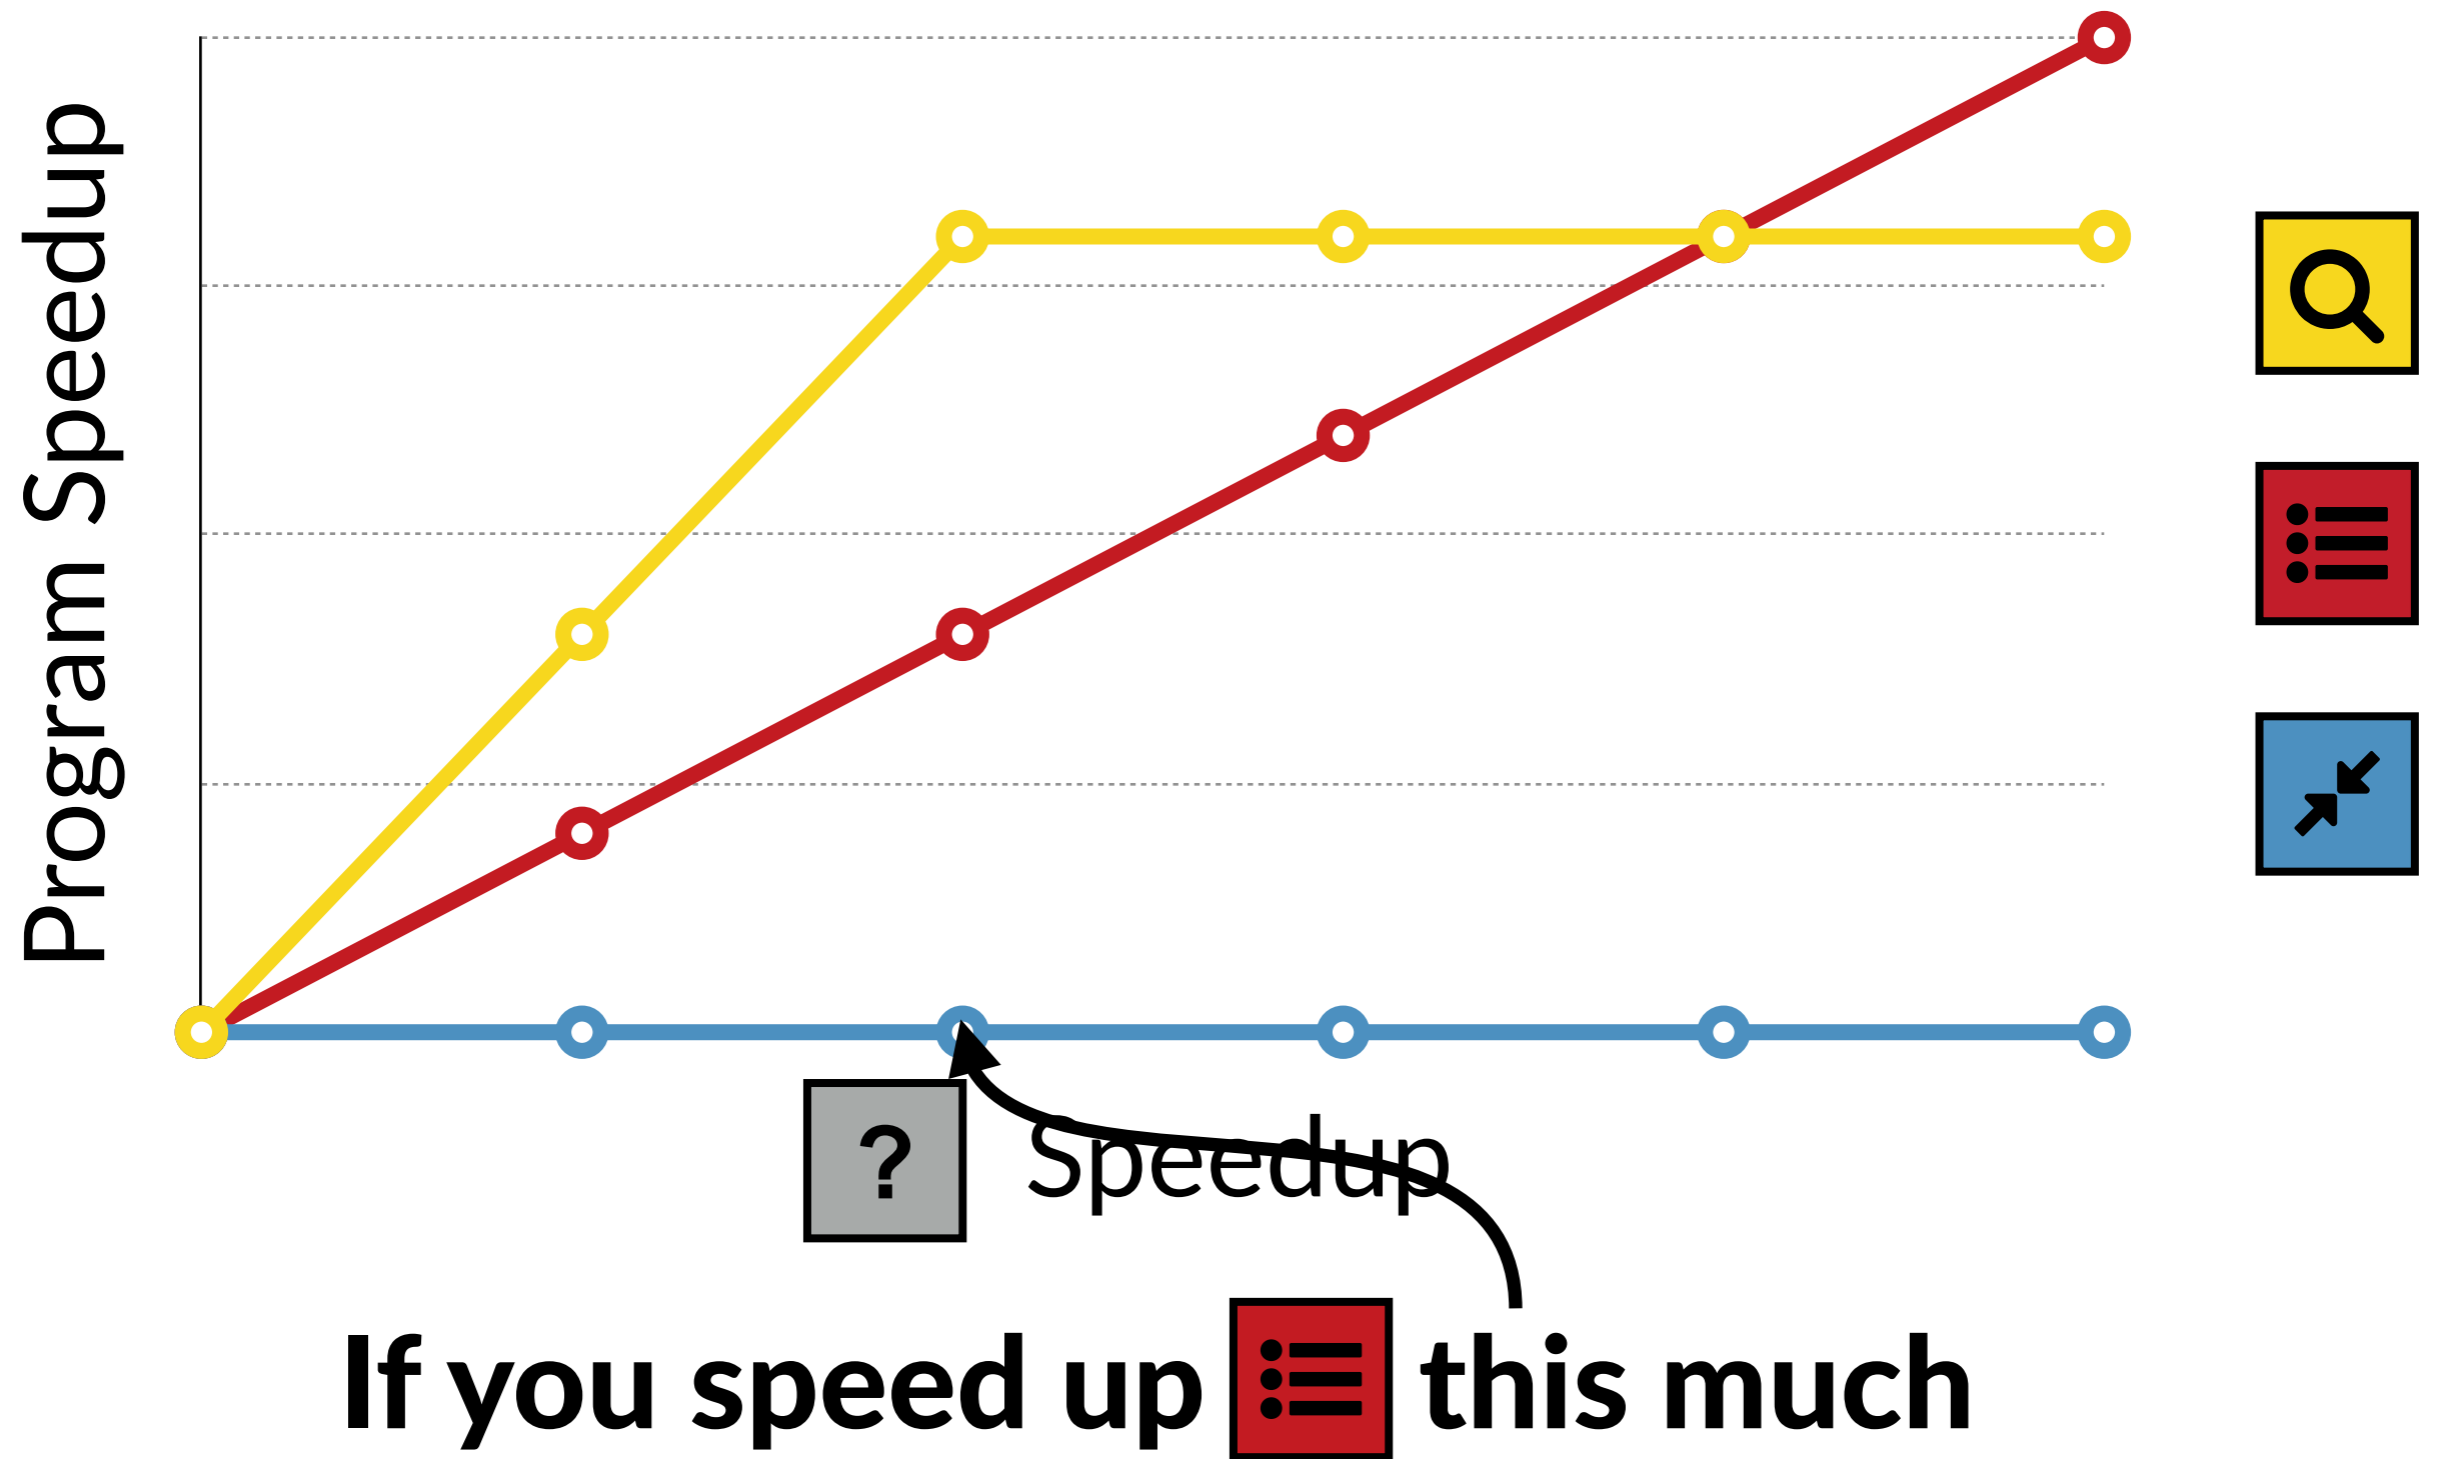
Tells you where optimizations will make a difference



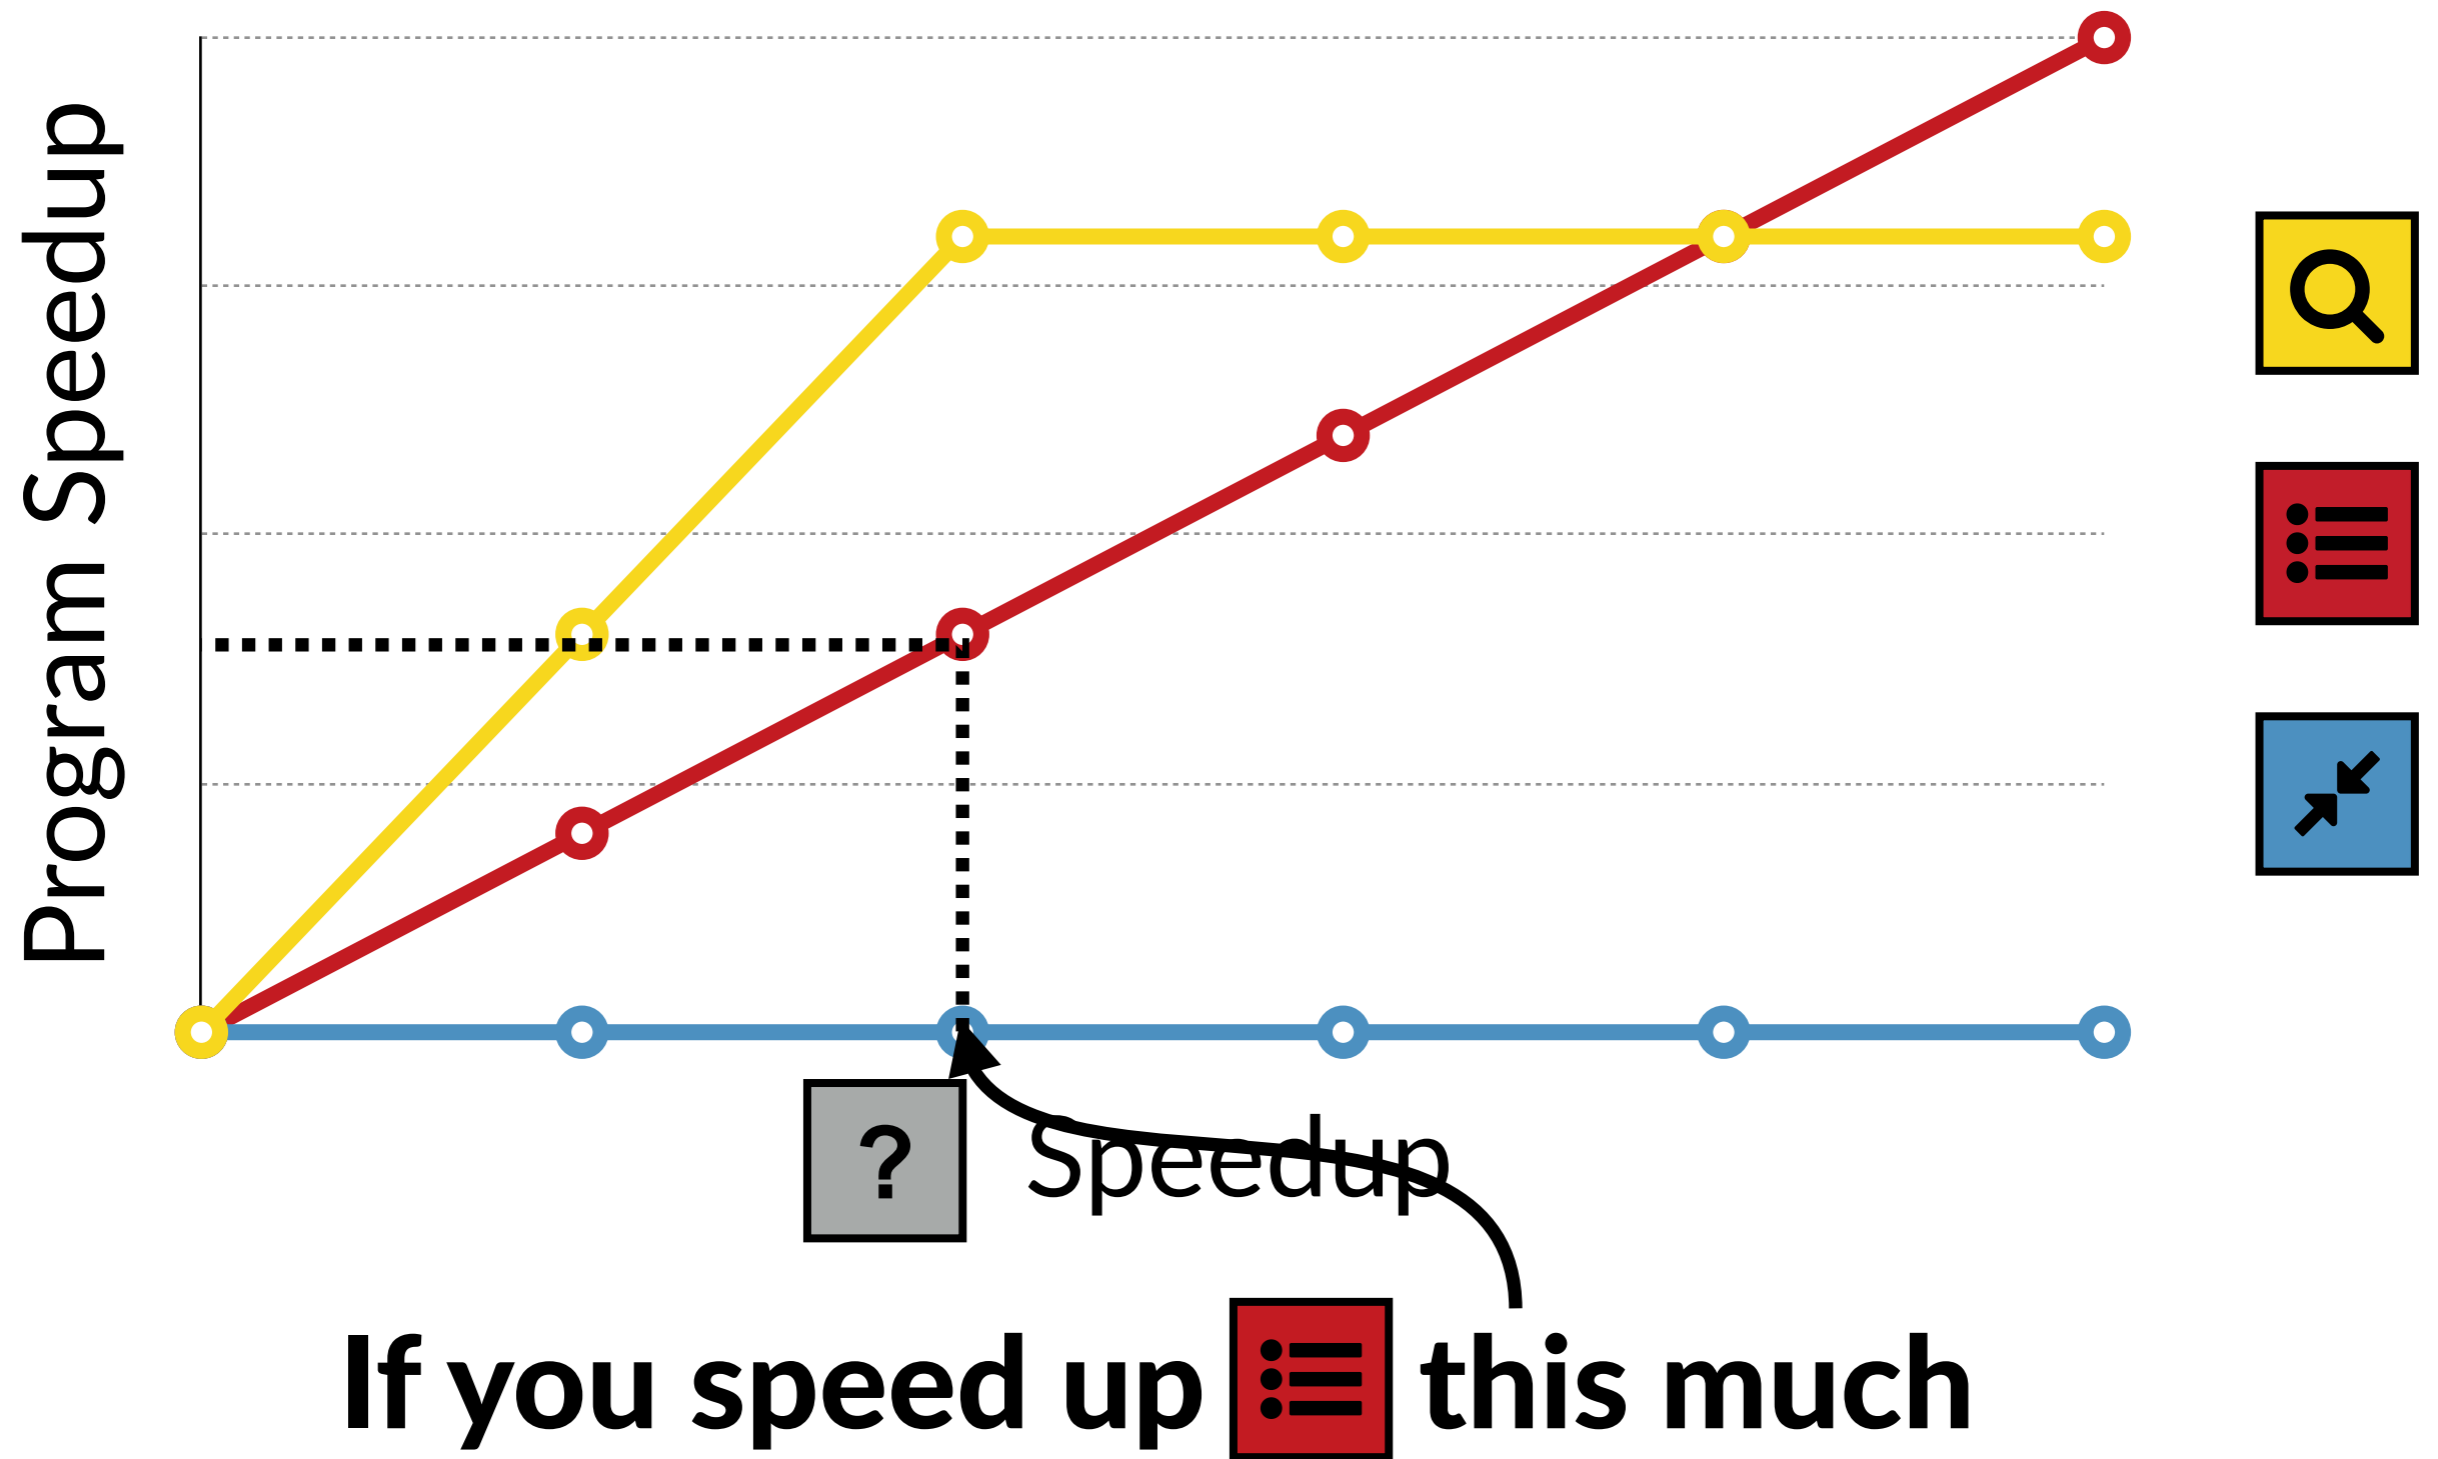
Causal Profile




Causal Profile



Causal Profile




Causal Profile




If you speed up  this much

The program will run this much faster

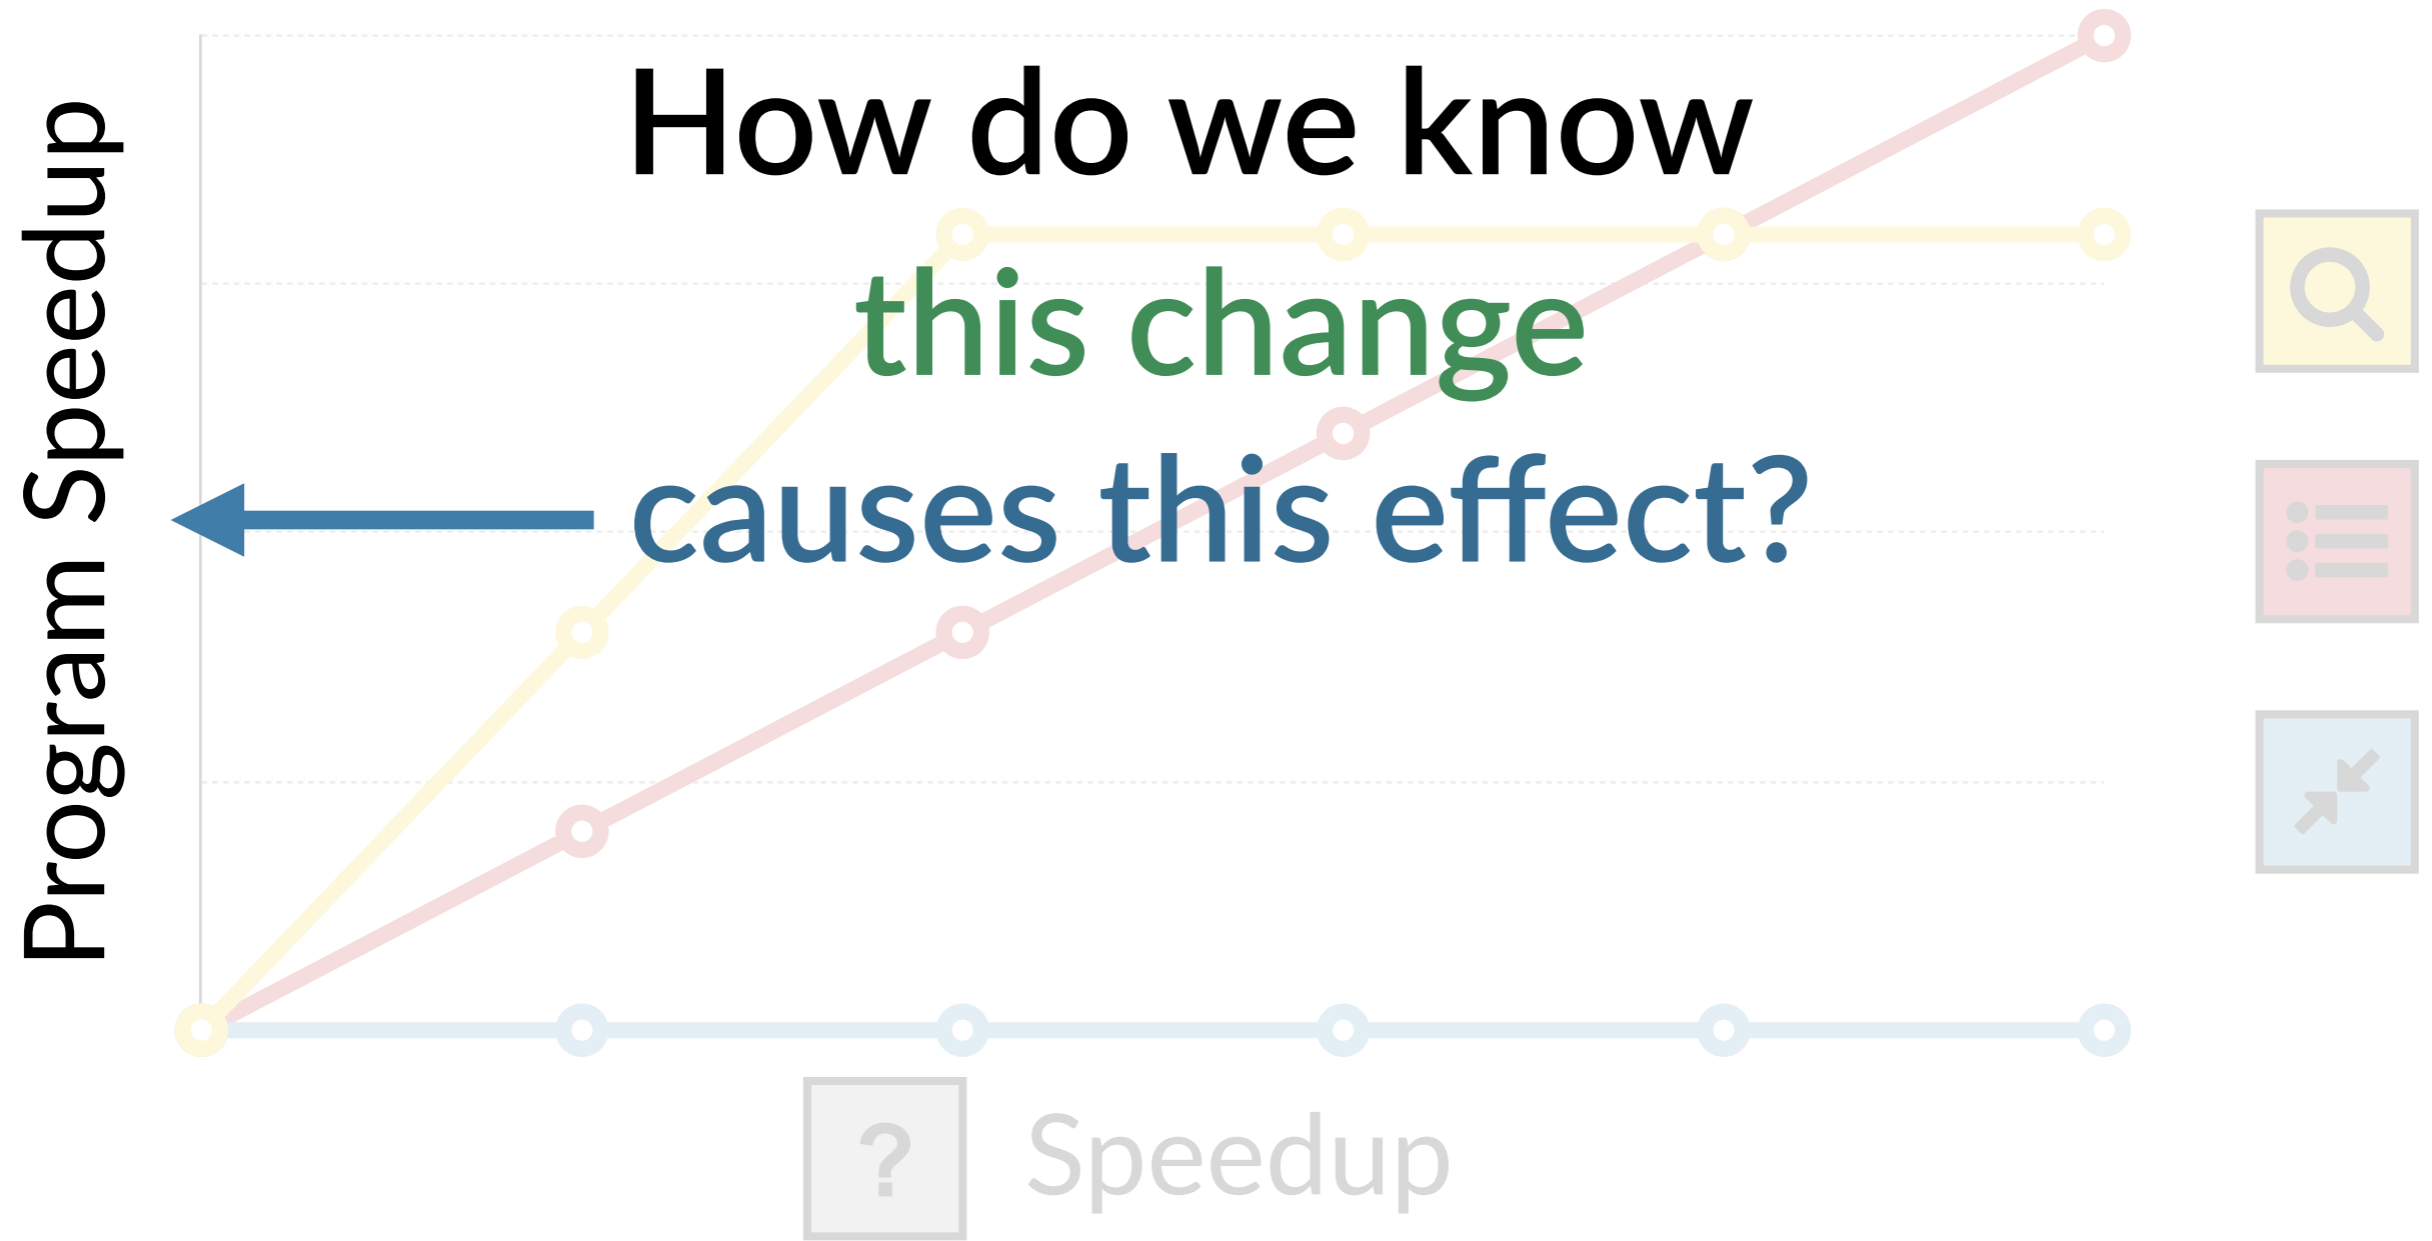
Causal Profile



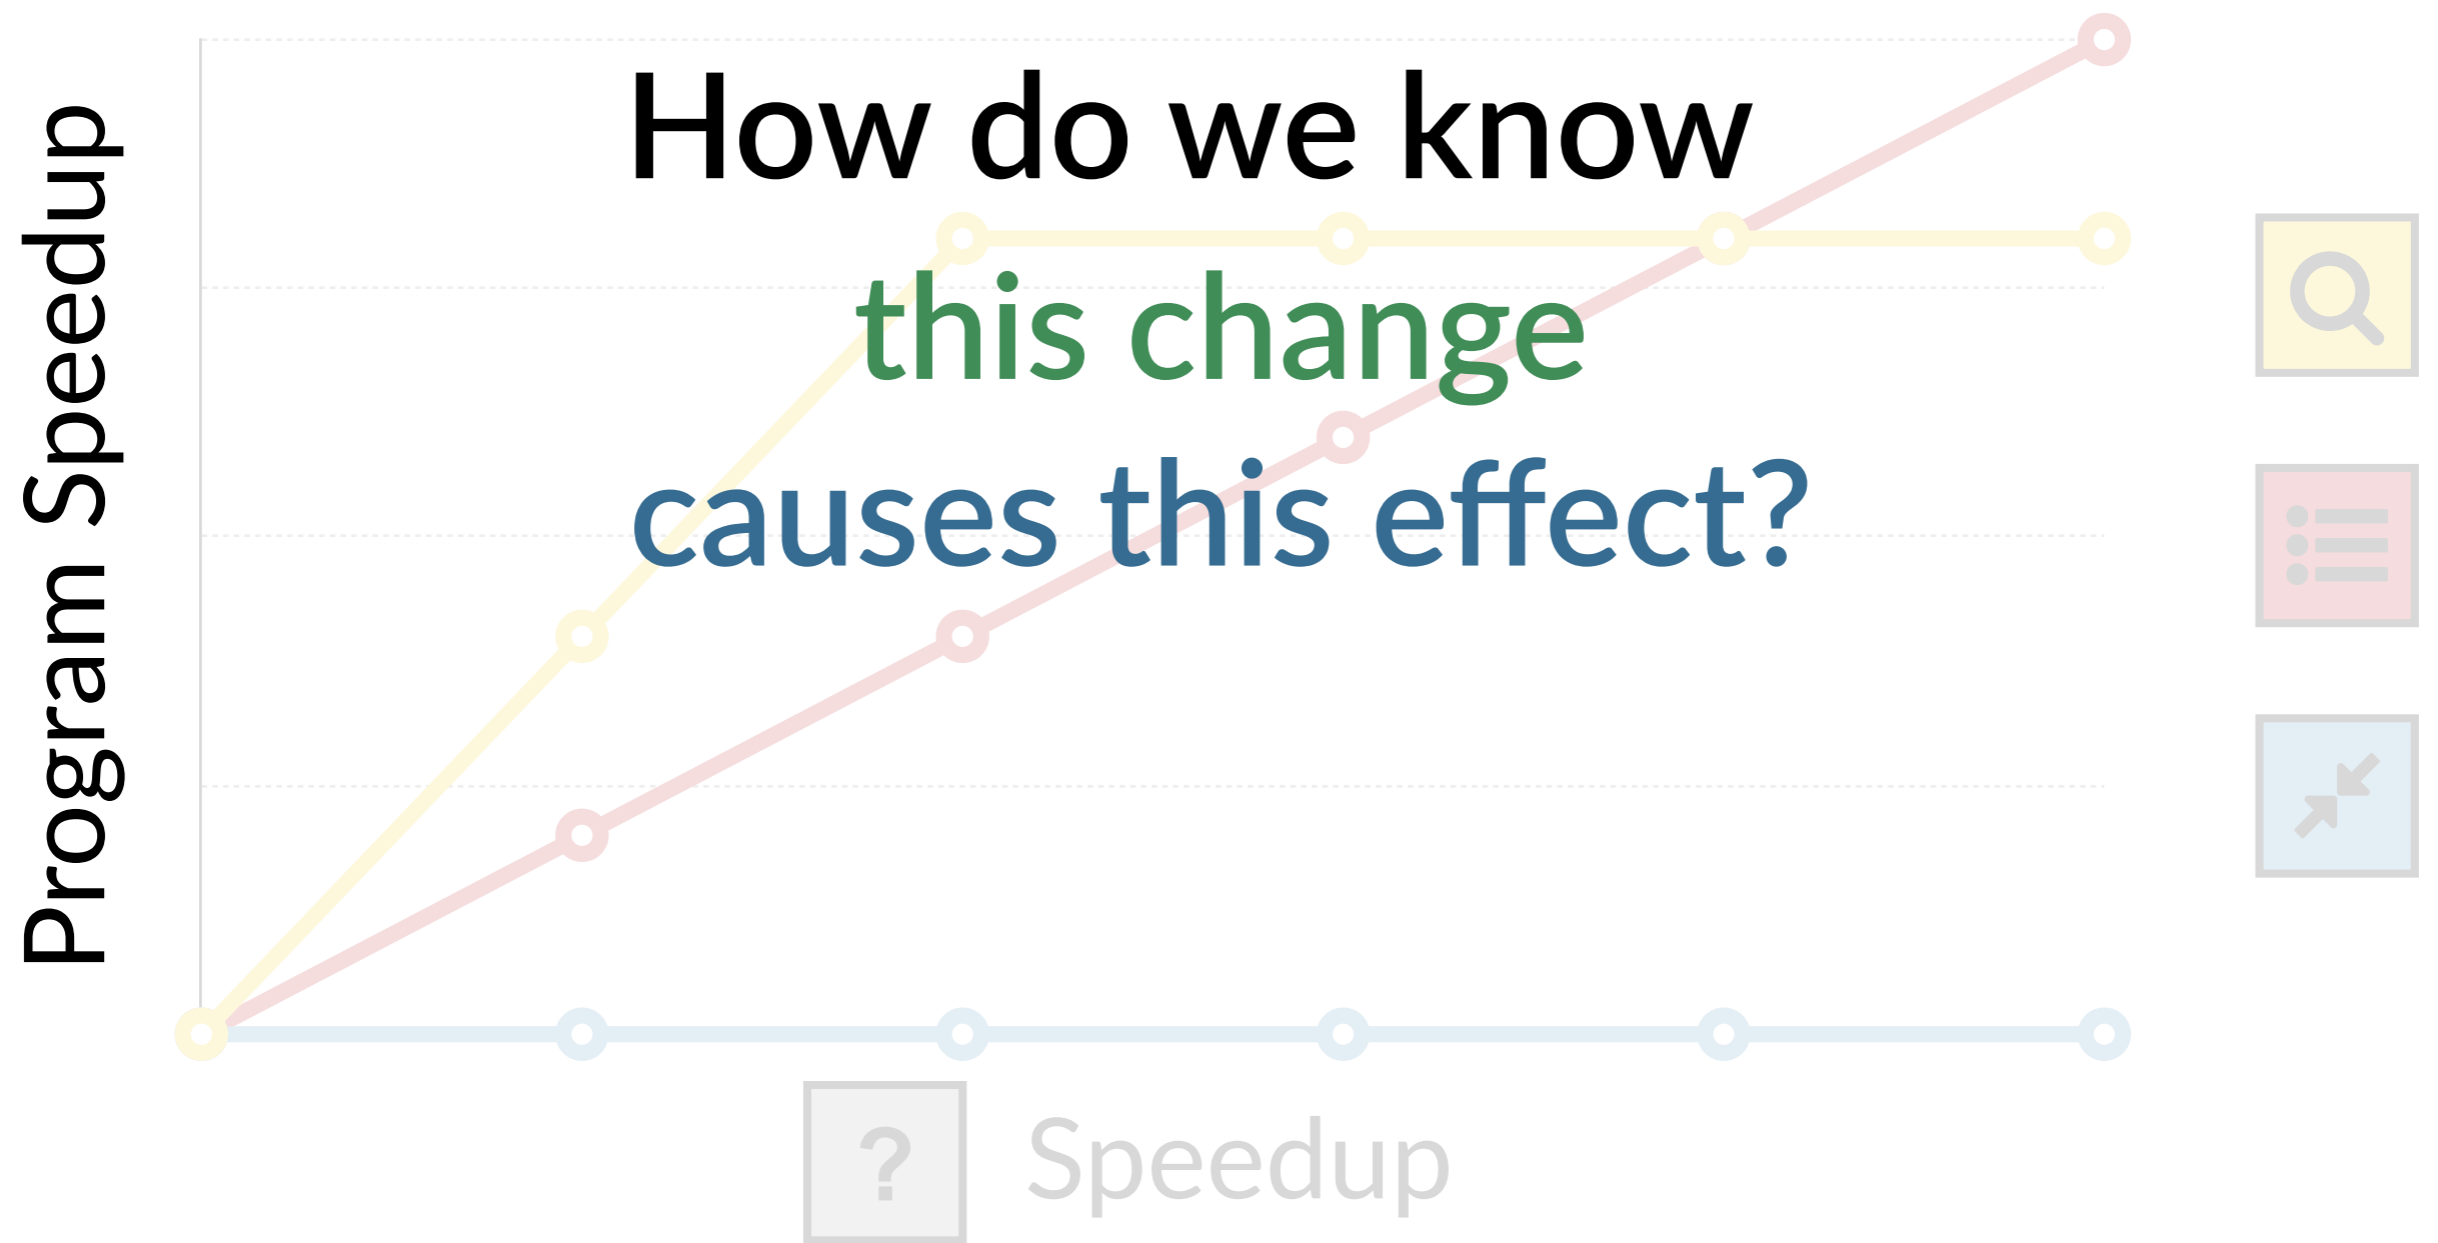
Causal Profile



Causal Profile




Causal Profile




Run an experiment

Performance Experiments




Performance Experiments

If we could magically speed up  ...




Performance Experiments

If we could magically speed up  ...



Performance Experiments


If we could magically speed up ...



Speeding up  by this much...

Performance Experiments


If we could magically speed up  ...



Speeding up  by this much...
speeds up the program by this much.


Performance Experiments

If we could magically speed up  ...



Performance Experiments


If we could magically speed up **Q**...



More speedup in **Q**...

Performance Experiments

If we could magically speed up **Q** ...




More speedup in **Q** ...

leads to a larger program speedup.


Performance Experiments

If we could magically speed up  ...




Performance Experiments

If we could magically speed up  ...



Performance Experiments


If we could magically speed up  ...



No program speedup


Performance Experiments

We're going to have to do this without magic.




Performance Experiments

We're going to have to do this *without magic*...



Performance Experiments


We're going to have to do this *without magic*...



Otherwise, we'd just do this...

Performance Experiments


We're going to have to do this *without magic*...



Otherwise, we'd just do this...


Virtual Speedup

“Speed up”  by slowing everything else down.



Virtual Speedup


“Speed up”  by slowing everything else down.



Speeding up  by this much...


Virtual Speedup

“Speed up”  by slowing everything else down.


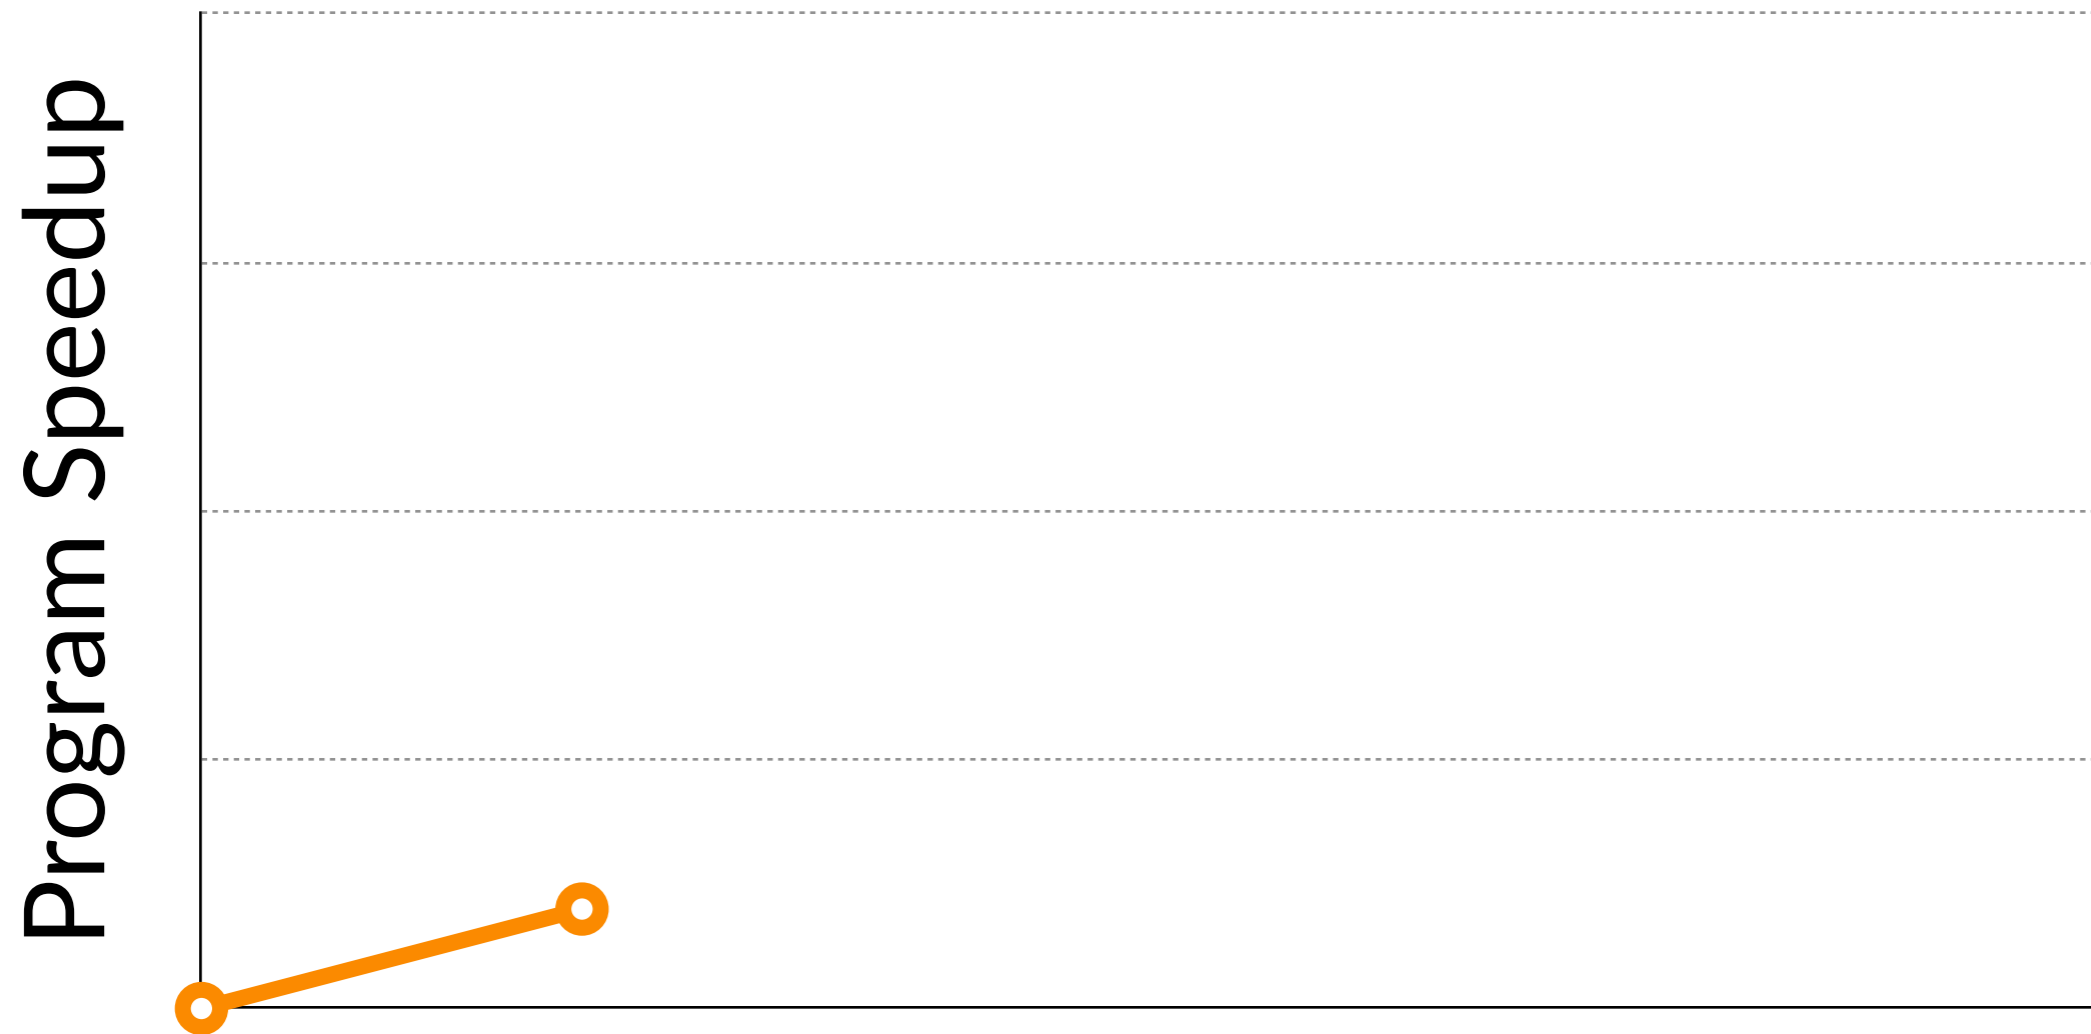


Speeding up  by this much...
speeds up the program by this much.

Speedup Results




Speedup Results



Speedup


Virtual Speedup

“Speed up”  by slowing everything else down.




Virtual Speedup

“Speed up”  by slowing everything else down.





Virtual Speedup

“Speed up”  by slowing everything else down.





A larger speedup has no additional effect

Speedup Results



Speedup


Speedup Results



Speedup

Virtual Speedup


“Speed up”  by slowing everything else down.



Each time  runs, pause all other threads.

Virtual Speedup


“Speed up”  by slowing everything else down.



Each time  runs, pause all other threads.



Virtual Speedup

“Speed up”  by slowing everything else down.




Each time  runs, pause all other threads.

Speedup Results




Speedup


Speedup Results



Speedup Results



Speedup Results










Take a picture



Add it to the database

Send it to Ogle

Find similar pictures

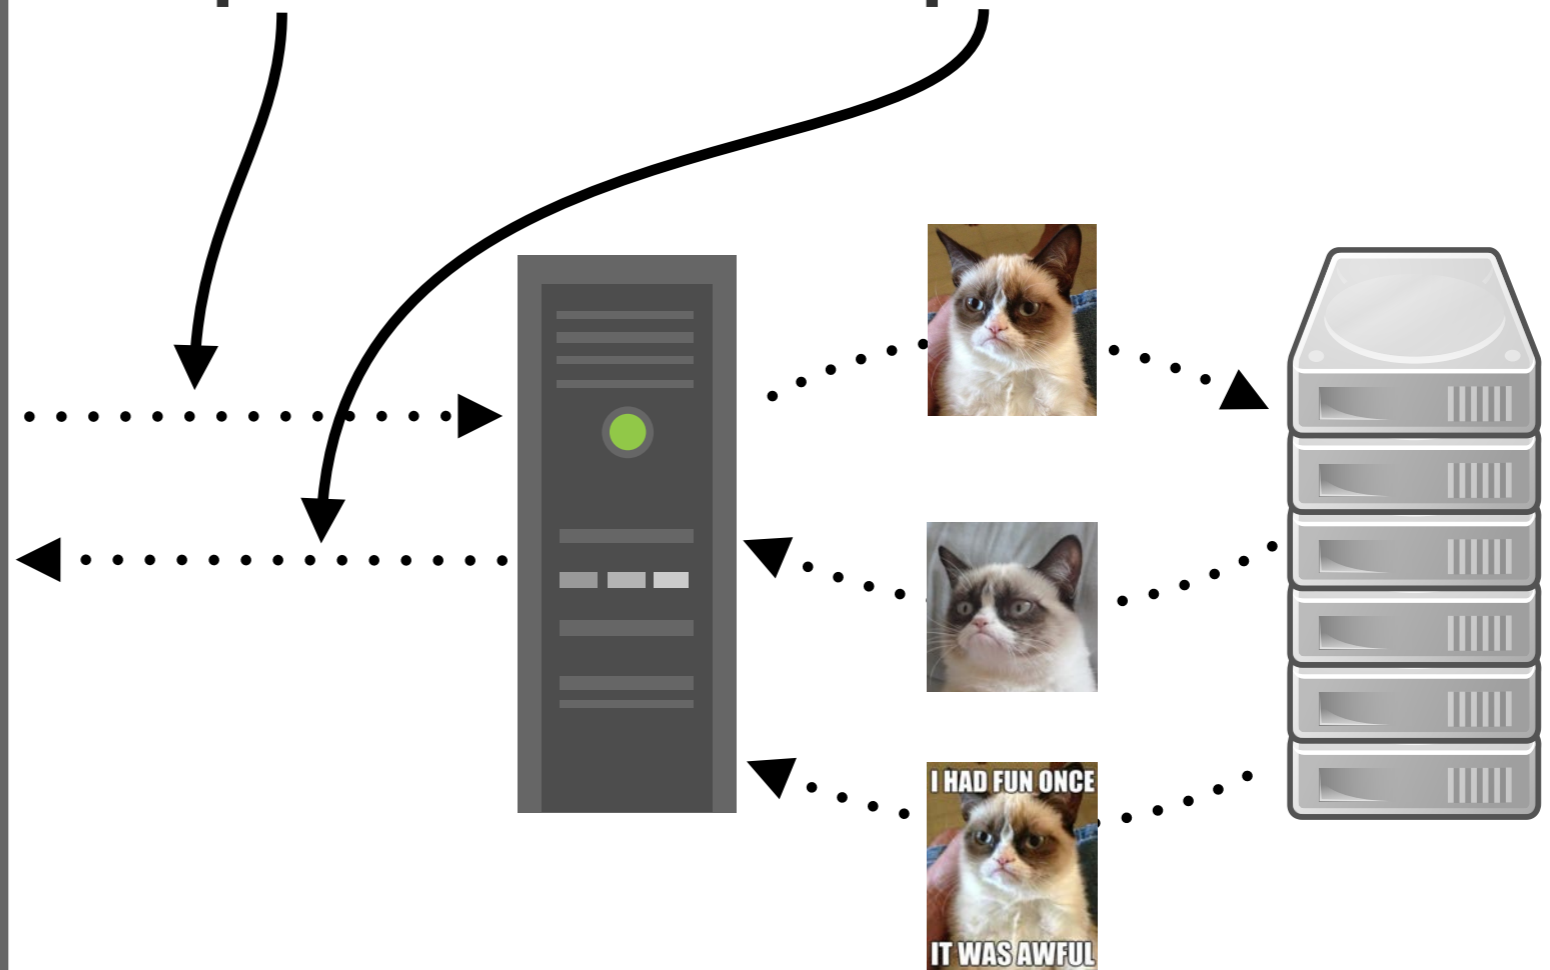

Send results







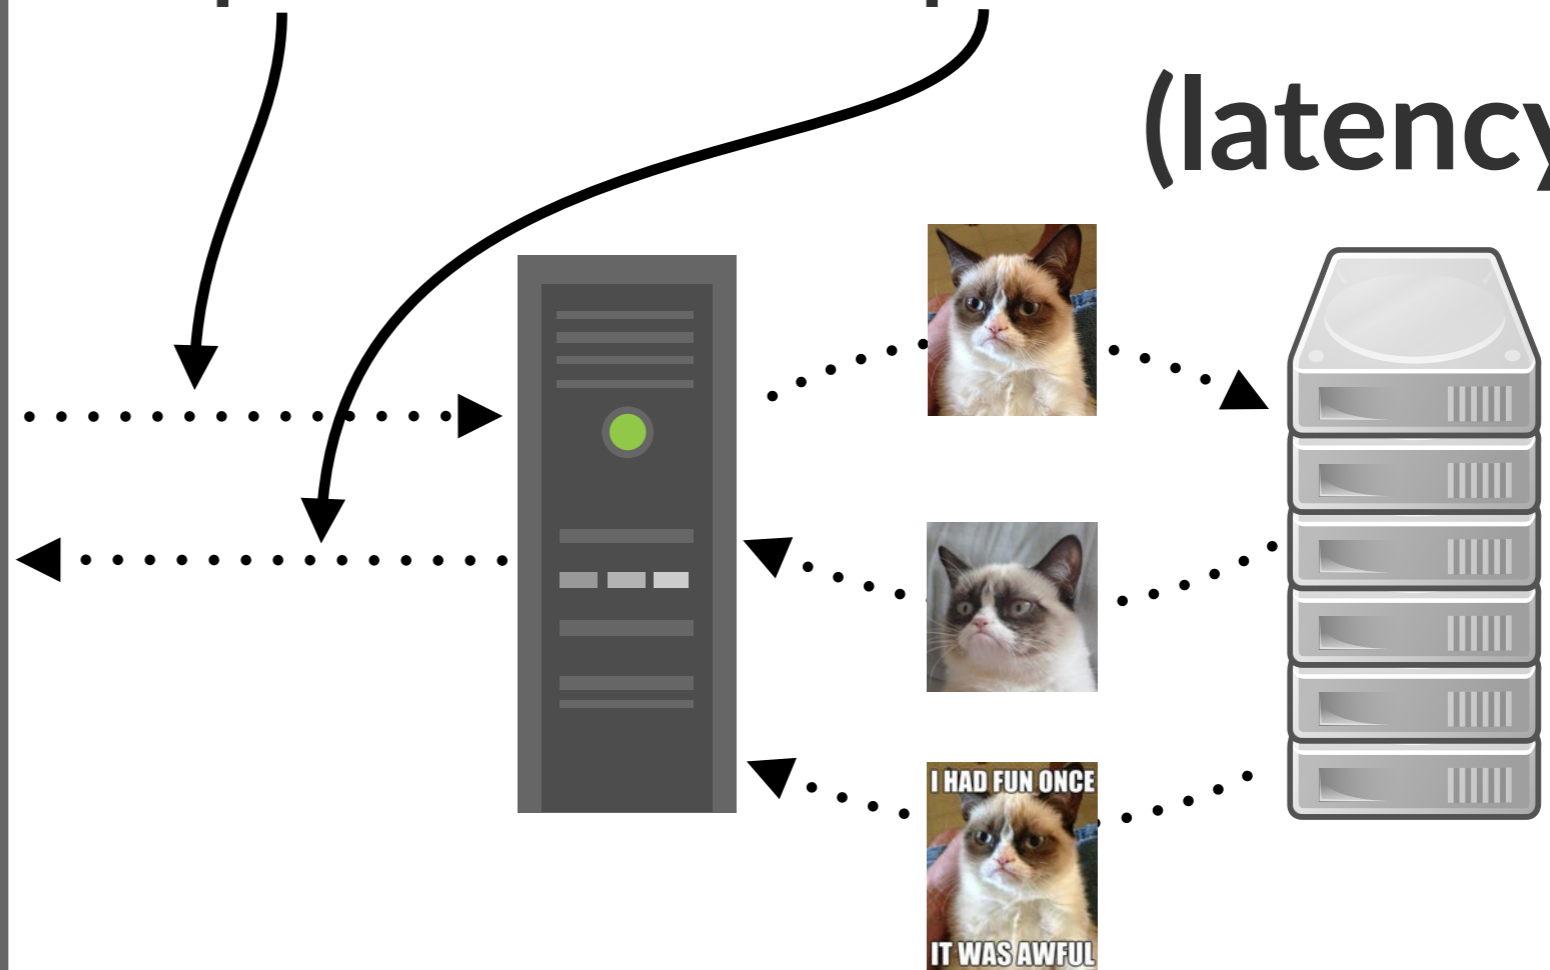
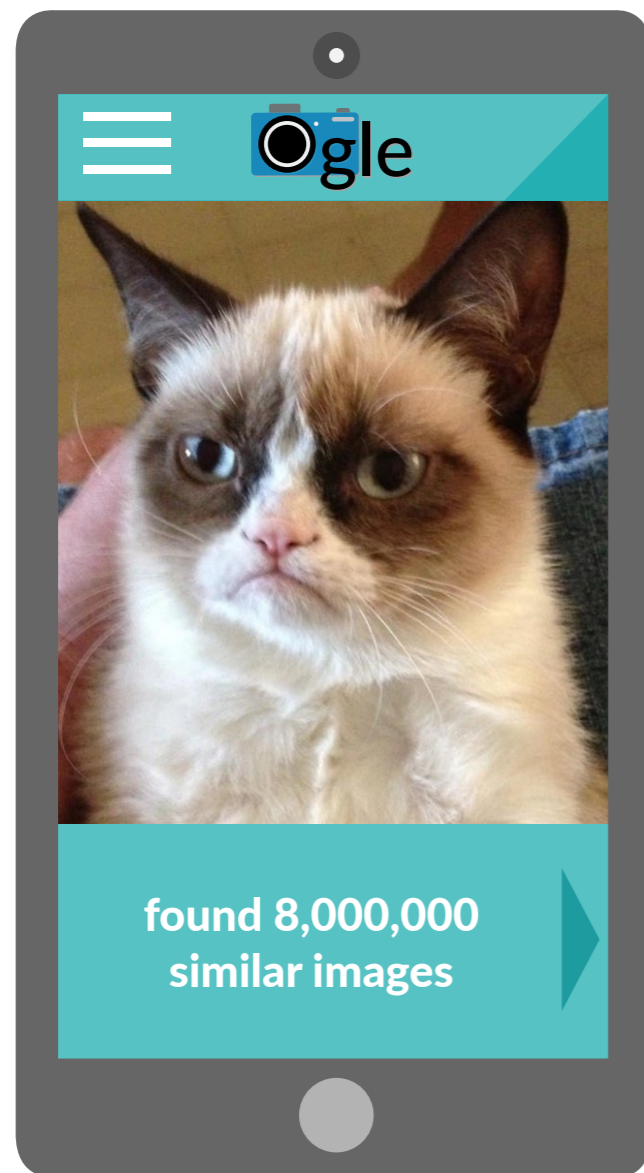
How long between request and response?





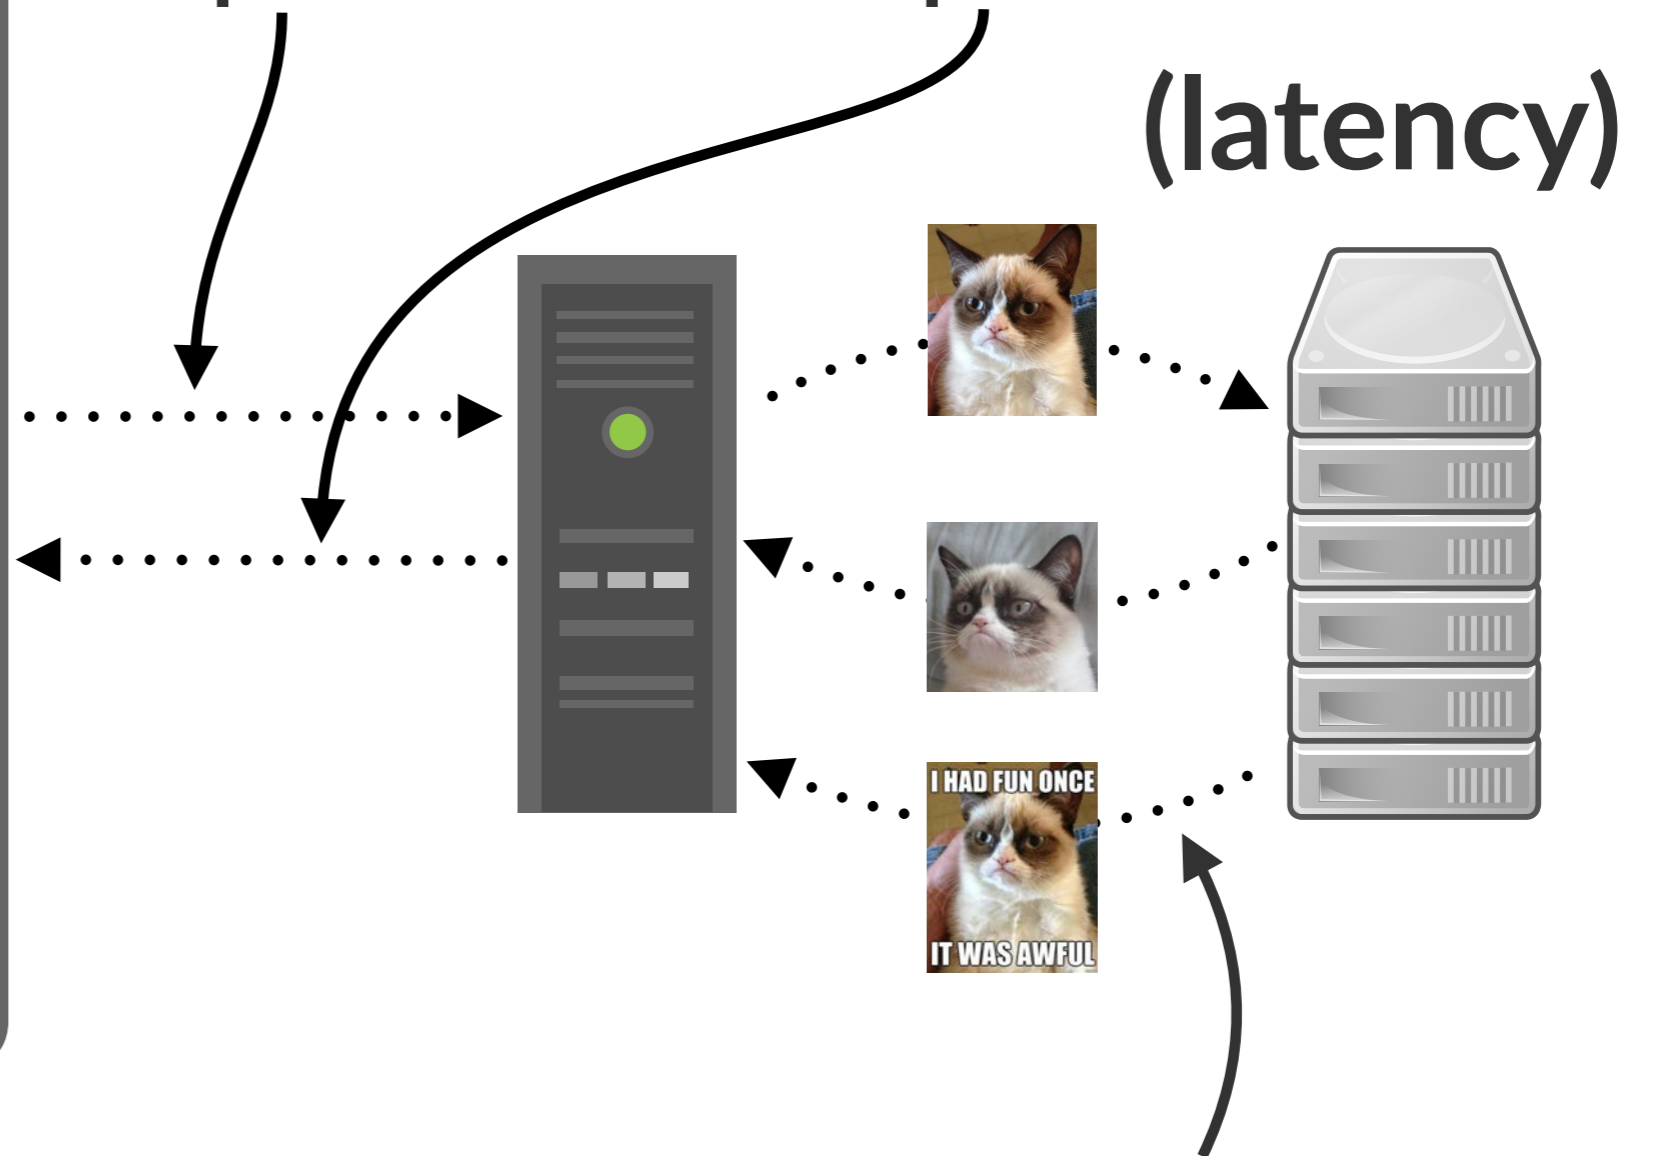

How long between request and response?

(latency)





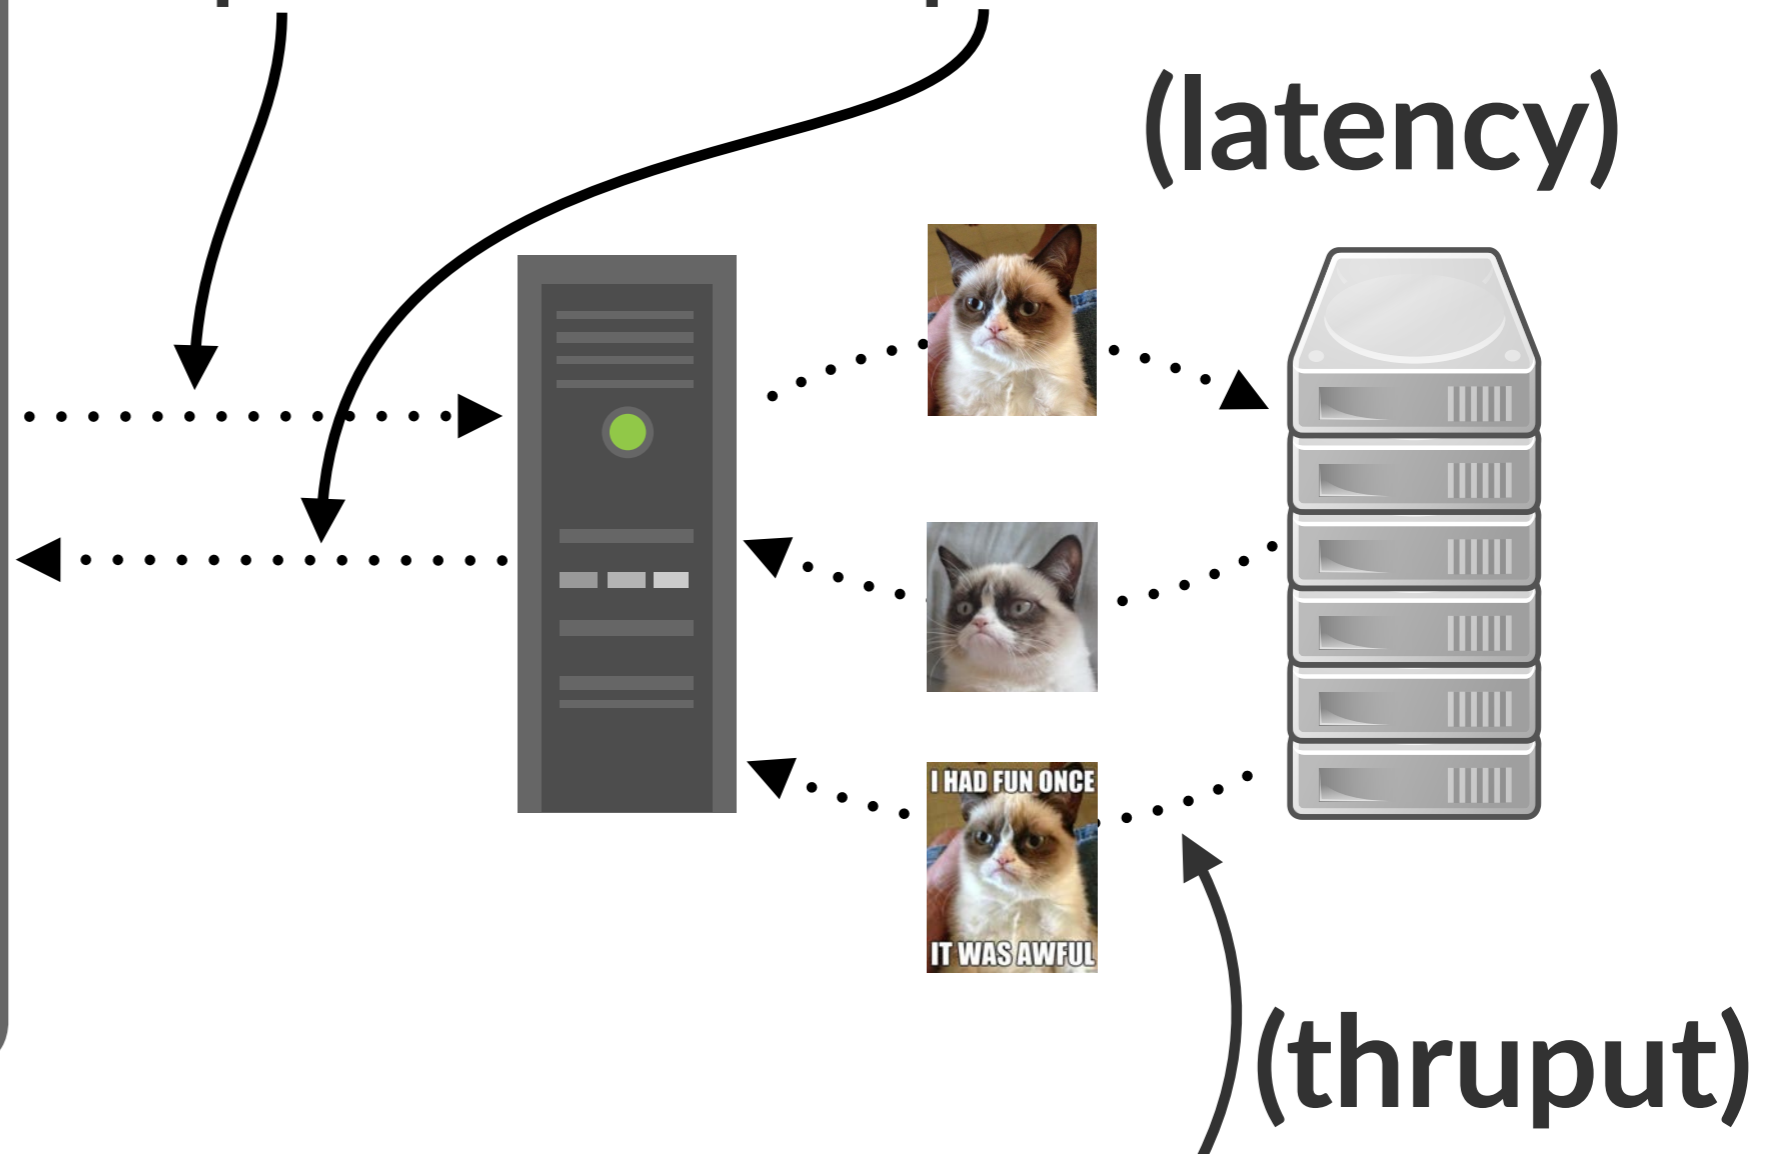

How long between request and response?



How fast do results come back?




How long between request and response?




How fast do results come back?

Progress Points

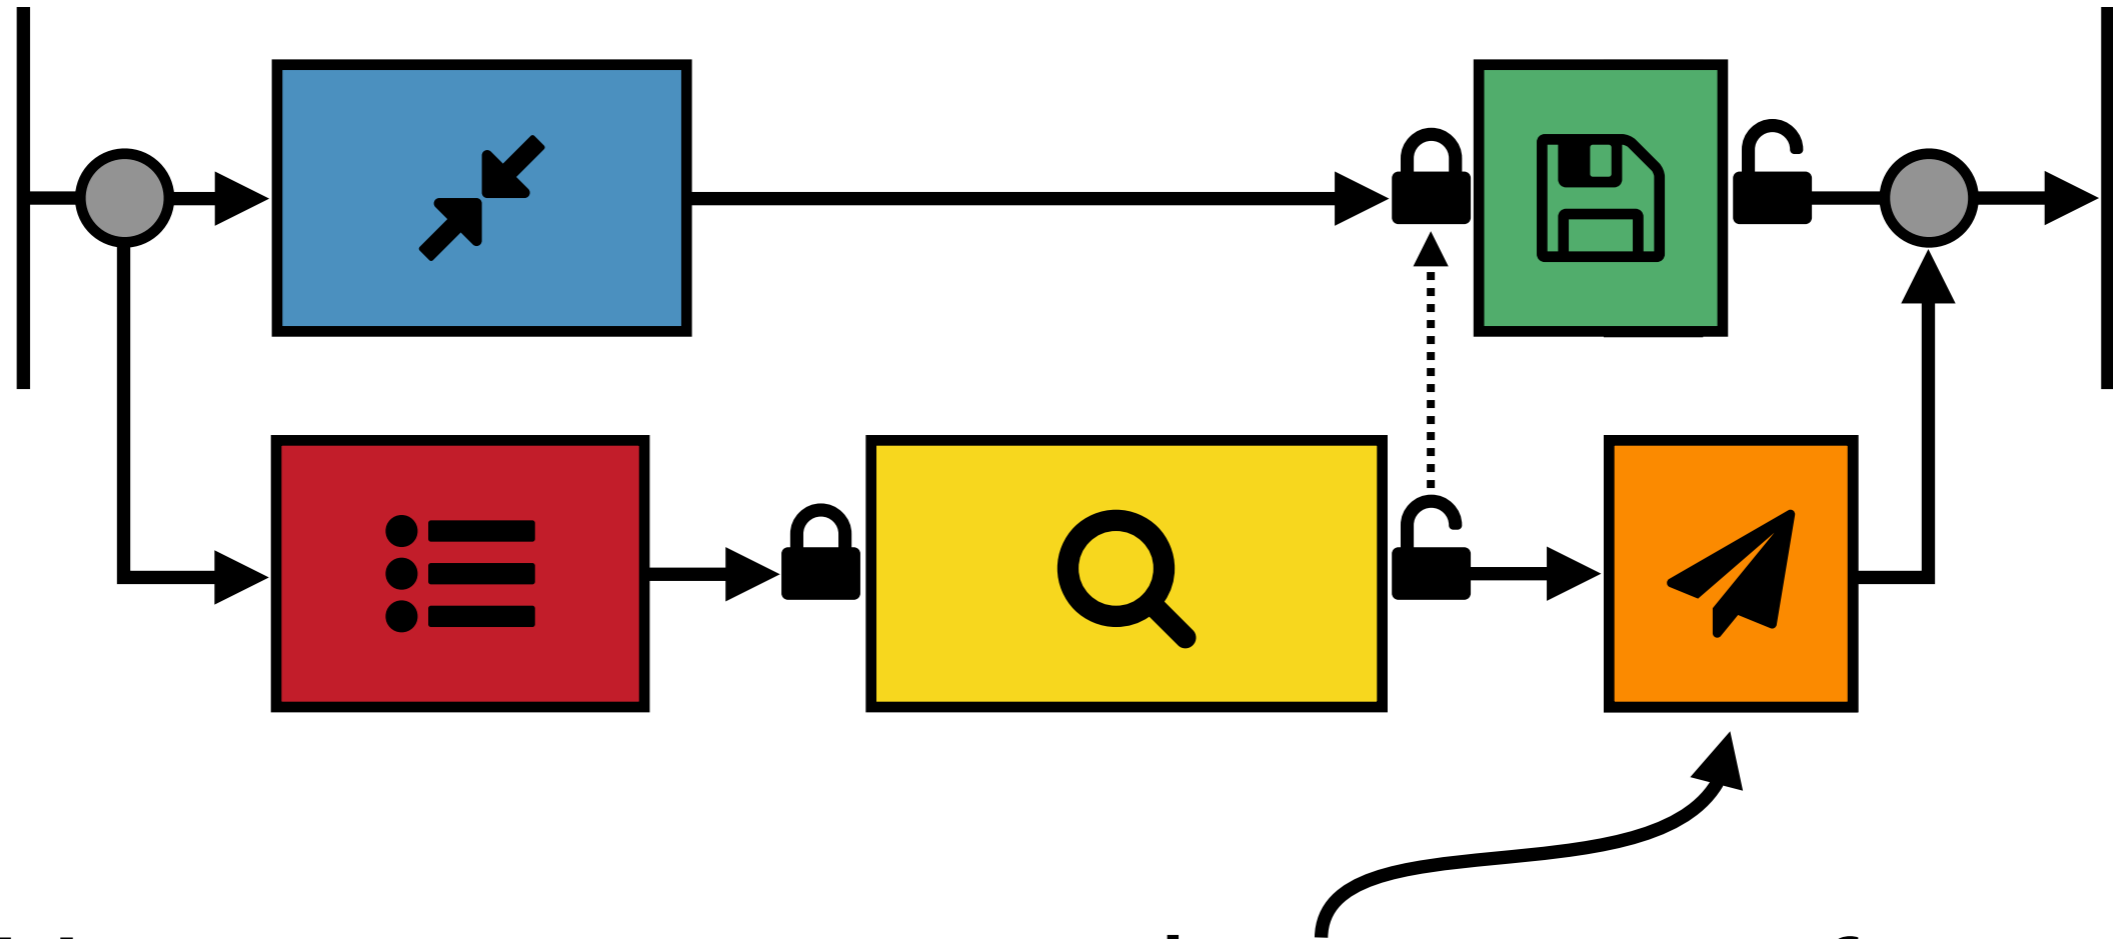


Progress Points



Homer wants to send responses faster.


Progress Points



Homer wants to send responses faster.

He marks  as a *progress point*.

Progress Points




Homer wants to send responses faster.

He marks  as a *progress point*.

Each time this code runs...

Progress Points

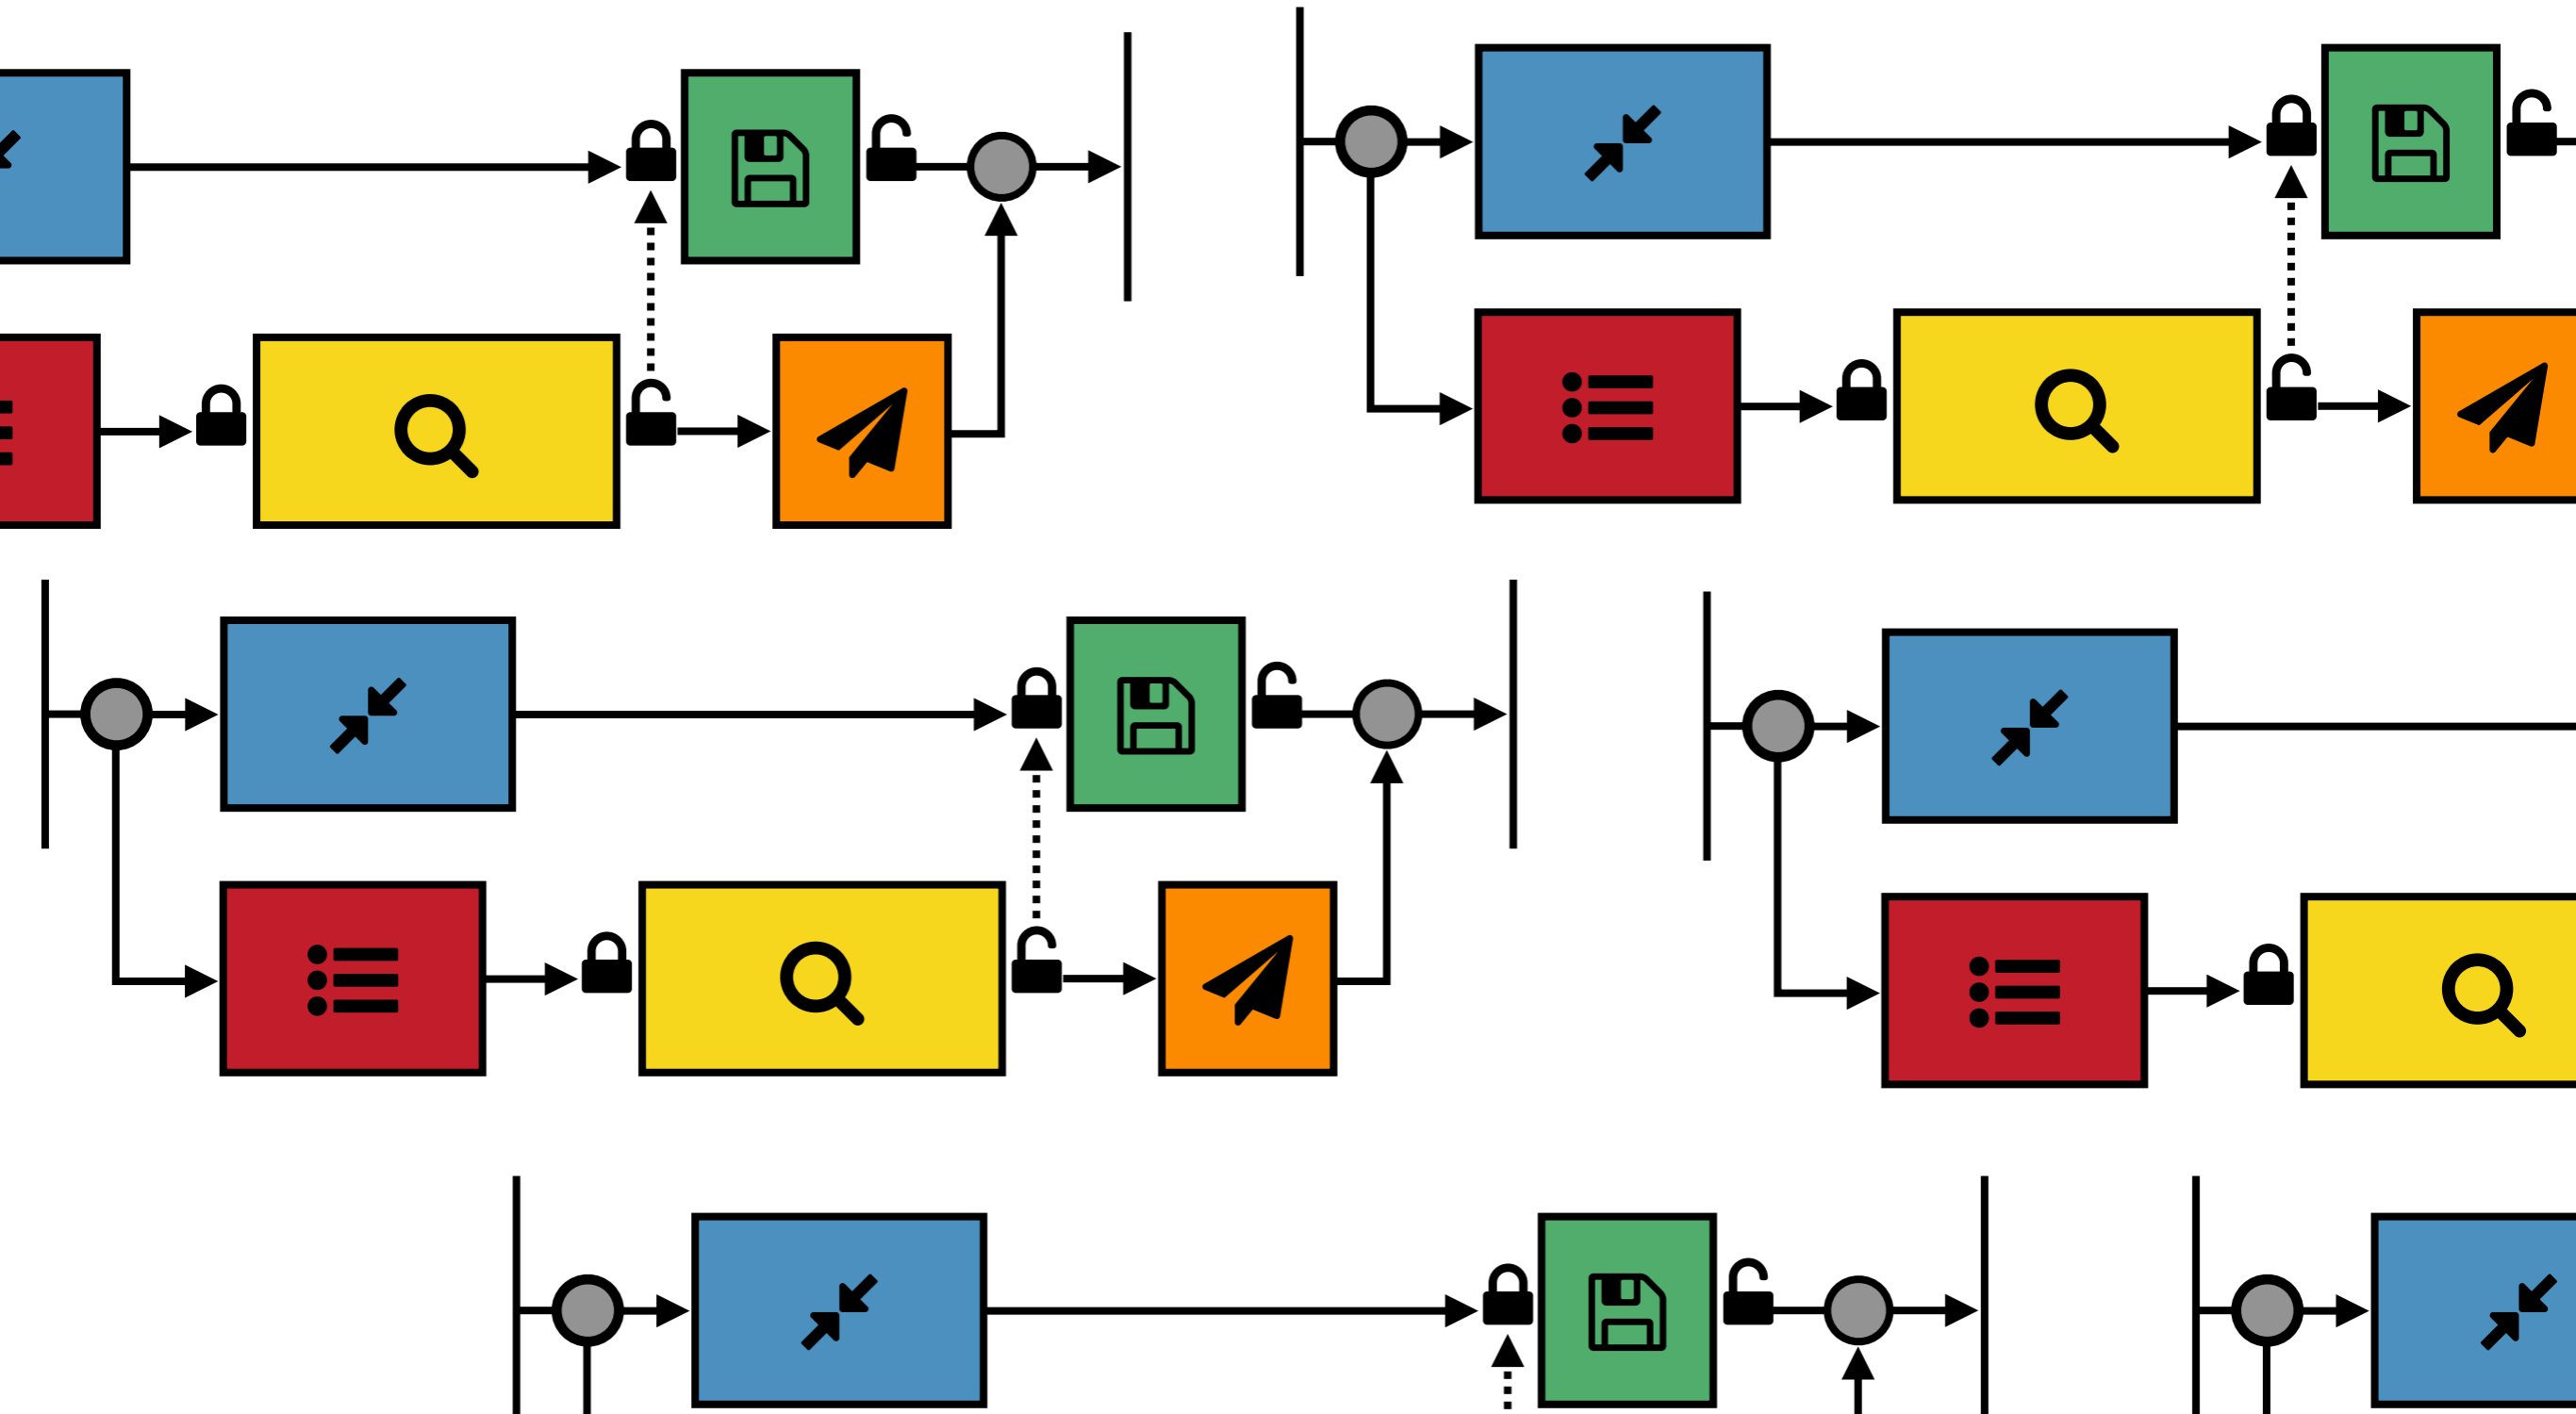


Homer wants to send responses faster.

He marks  as a *progress point*.


Each time this code runs...

Progress Points



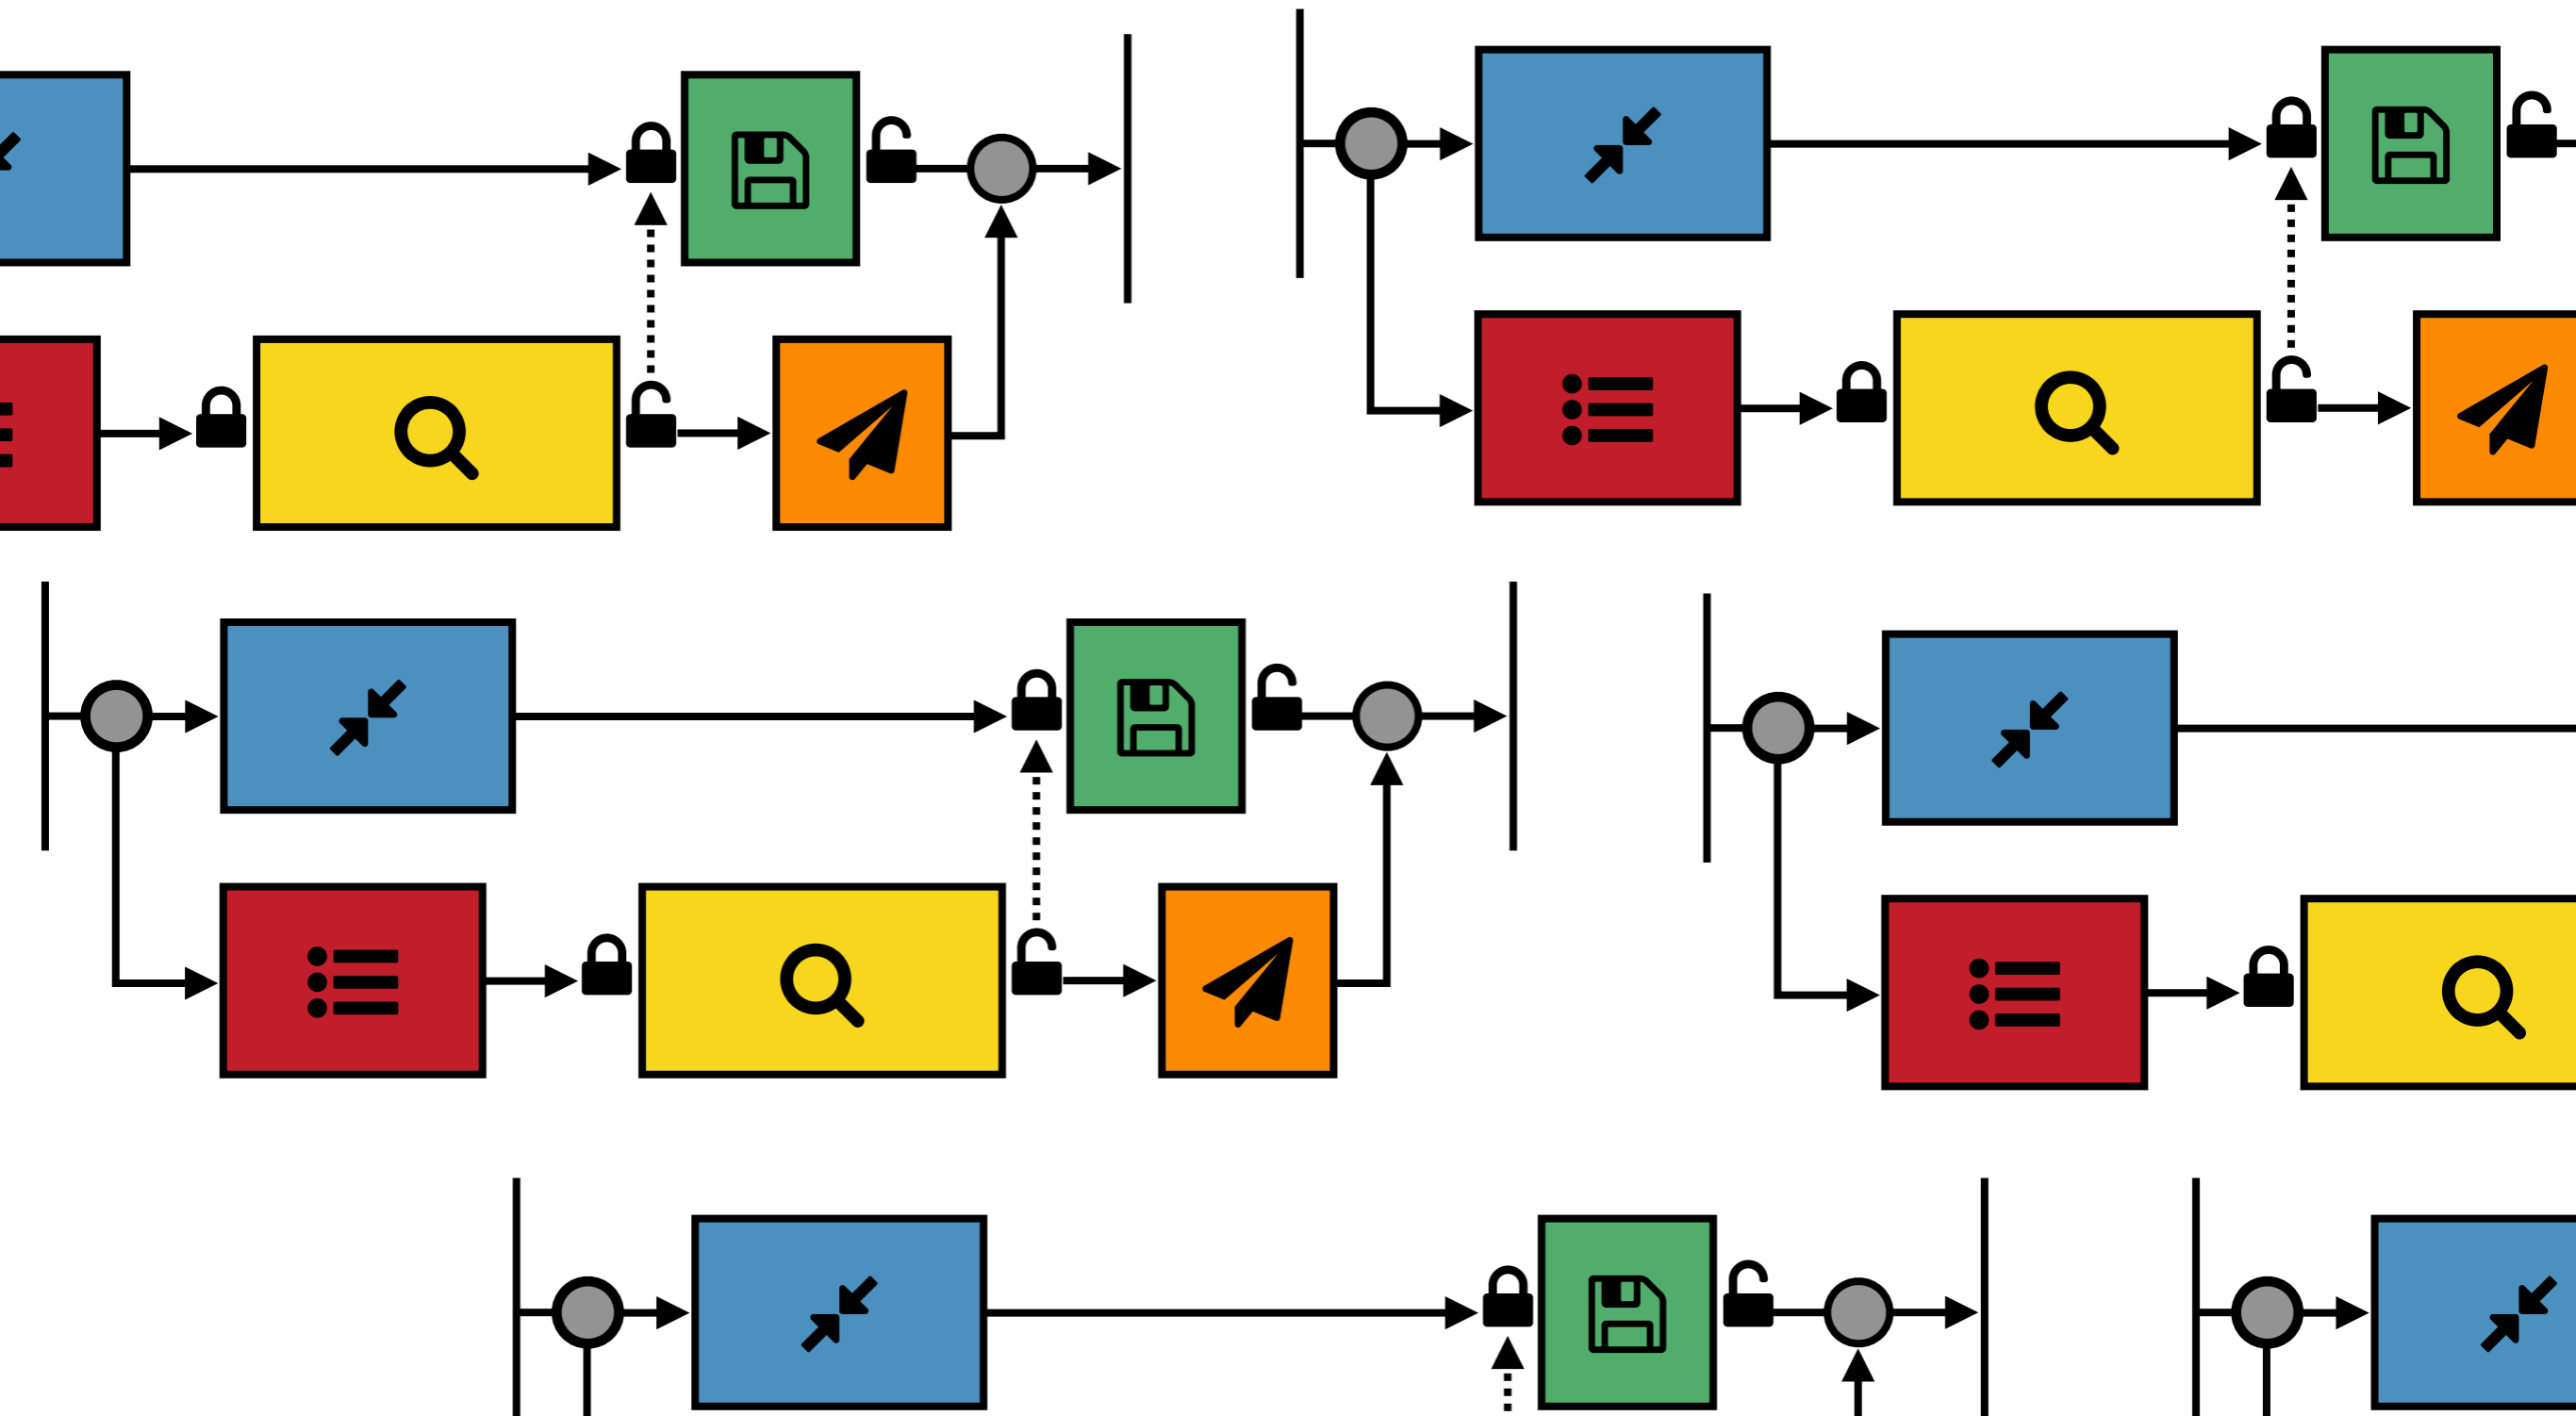
Progress Points

Many requests running for many users.

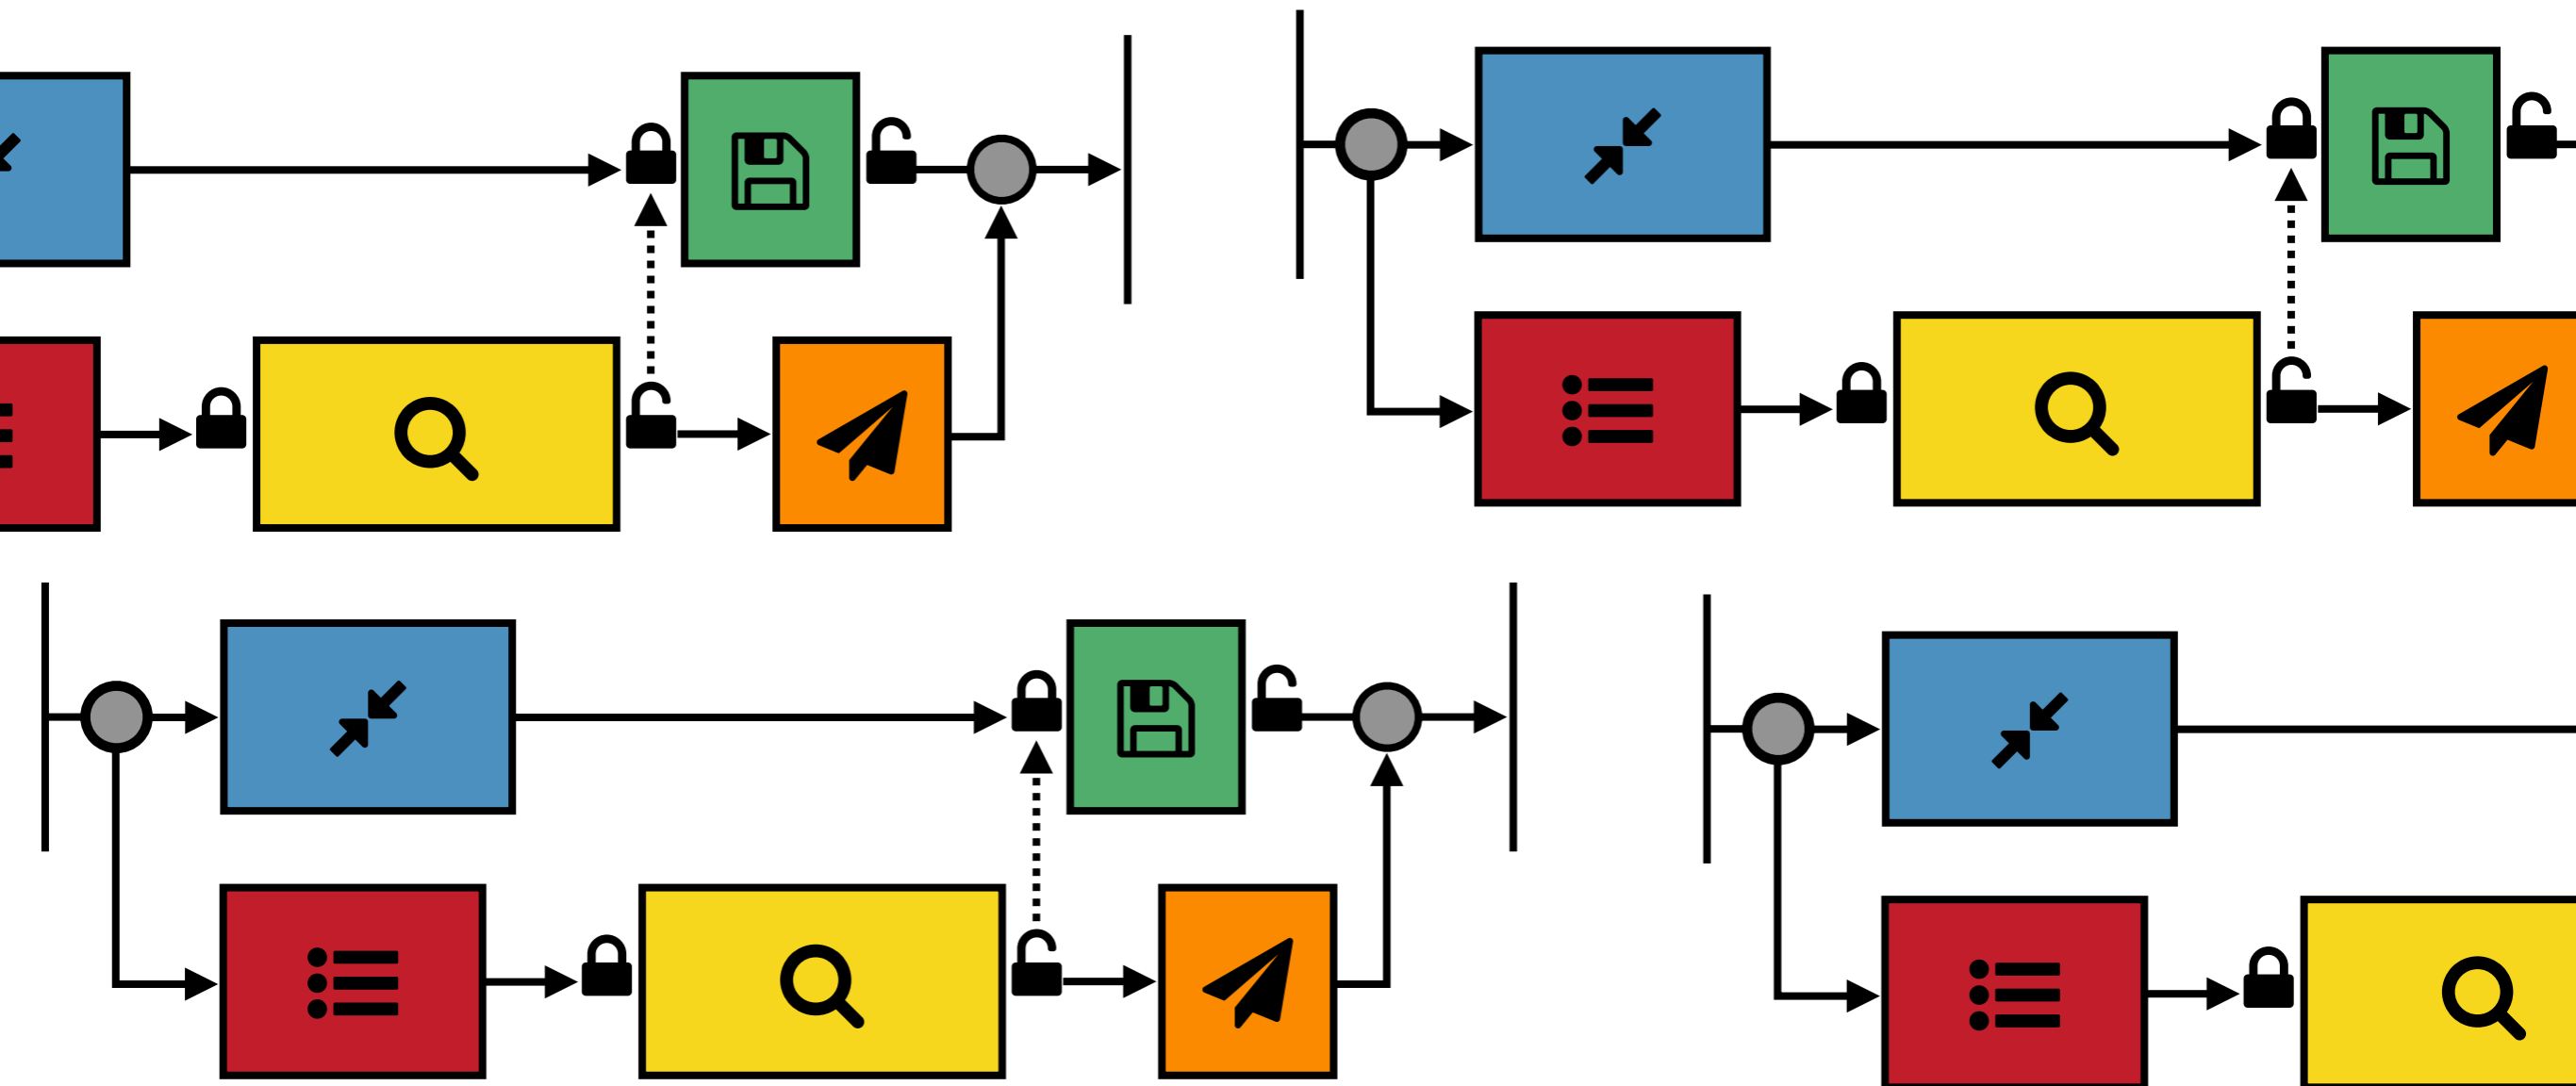


Progress Points

Many requests running for many users.




Progress Points



Progress Points

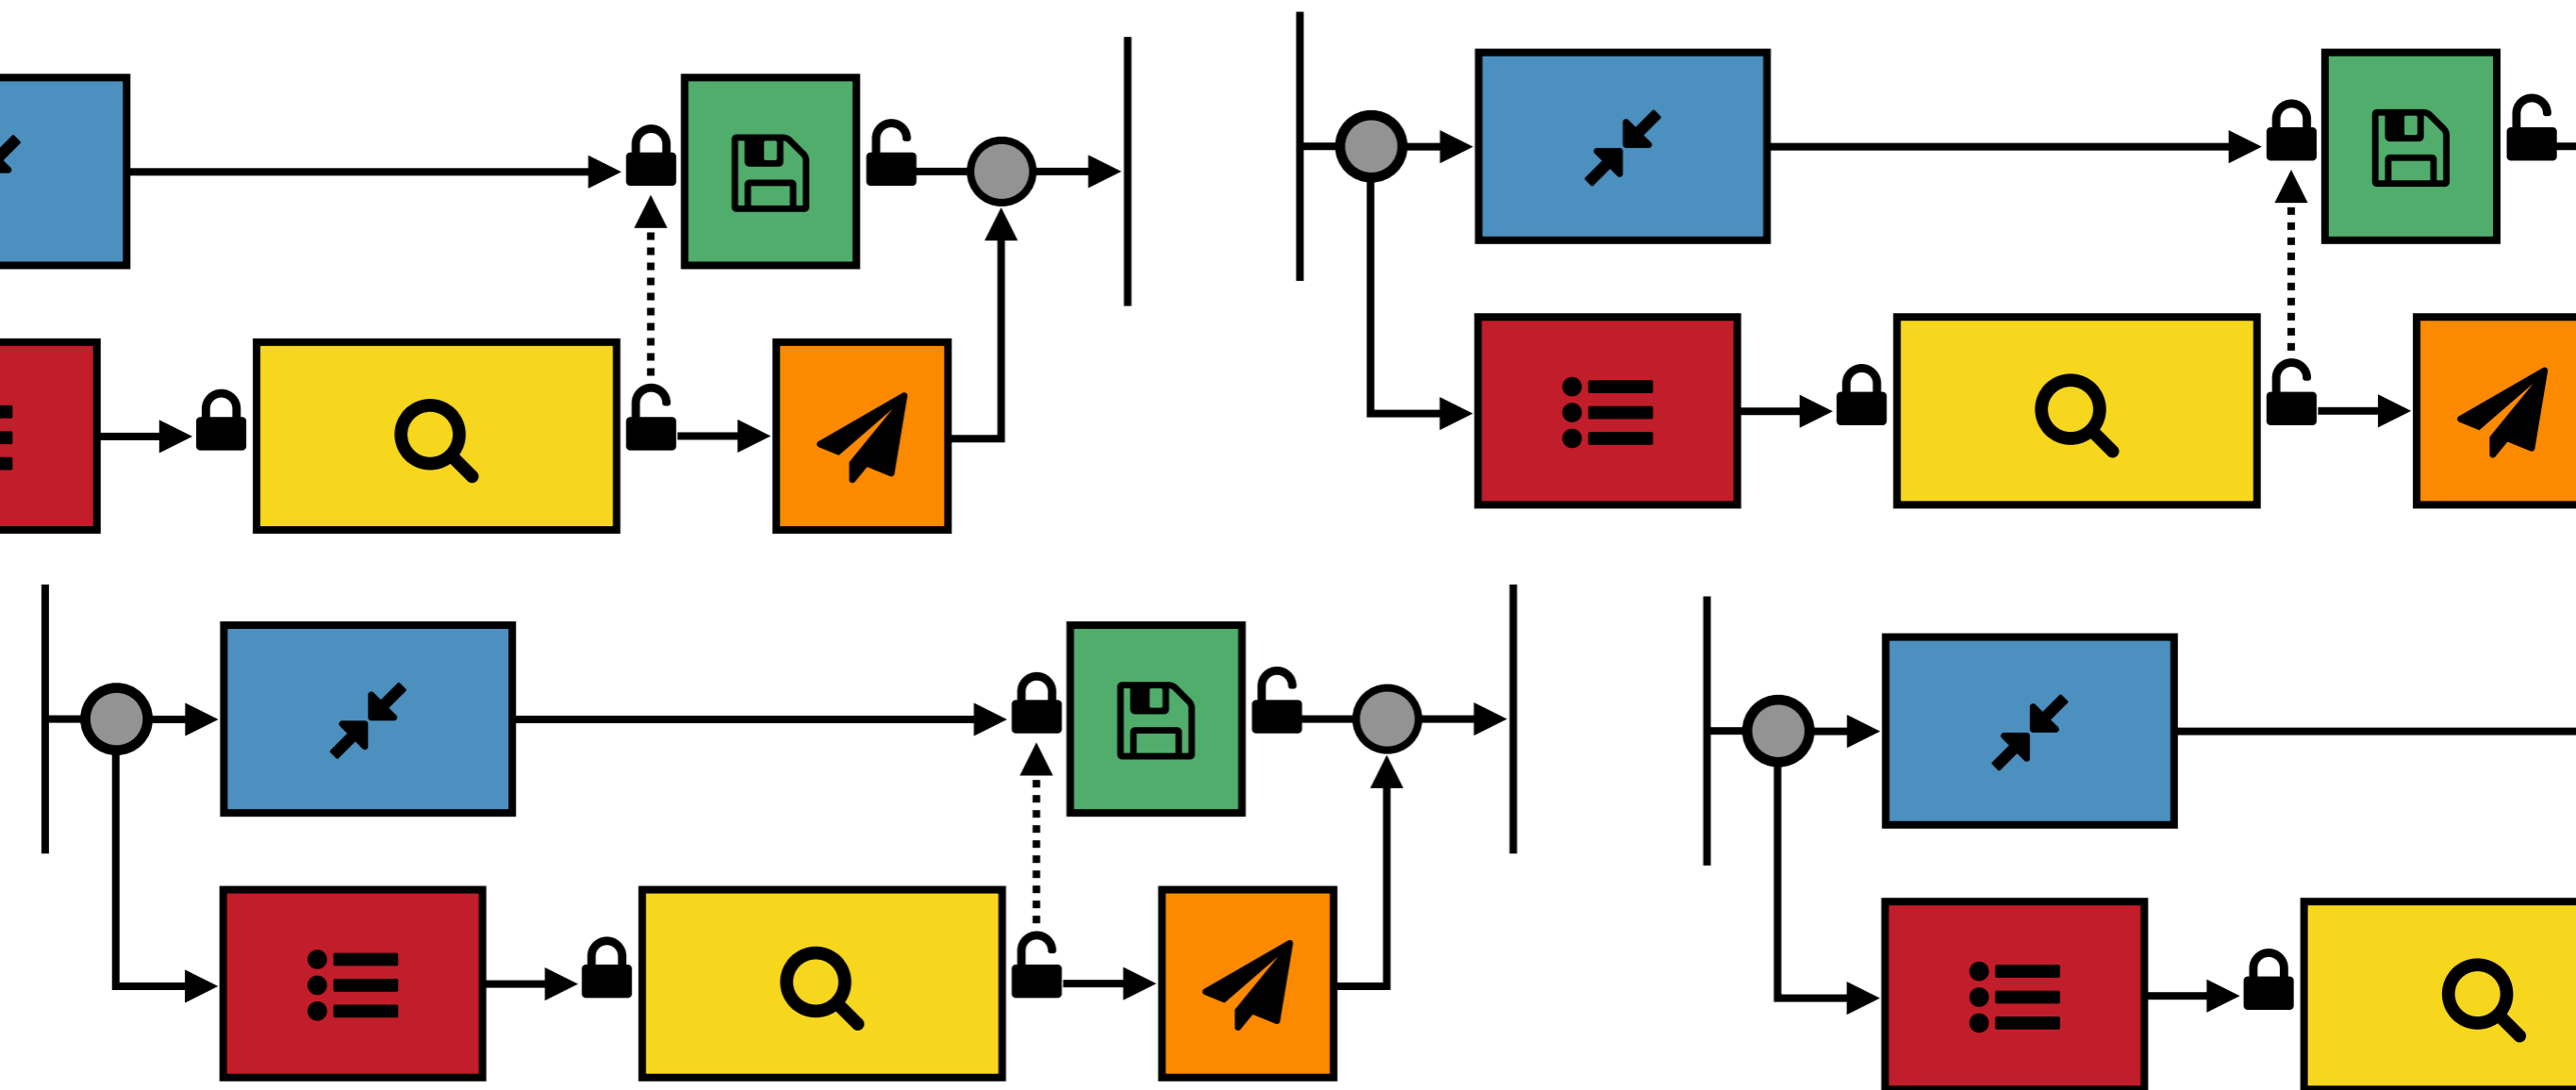
One progress point measures throughput.




Progress Points

One progress point measures throughput.


If I speed up , how much faster do I run ?



Progress Points




Progress Points



Homer wants to minimize response time.


Progress Points



Homer wants to minimize response time.

He adds *latency progress points*.


Progress Points



Homer wants to minimize response time.

He adds *latency progress points*.


Progress Points



Homer wants to minimize response time.

He adds *latency progress points*.


Progress Points



Homer wants to minimize response time.

Progress Points

Little's Law: $W = L / \lambda$




Homer wants to minimize response time.

Progress Points

Little's Law: $W = L / \lambda$

latency




Homer wants to minimize response time.

Progress Points

Little's Law: $W = L / \lambda$

latency = transactions




Homer wants to minimize response time.

Progress Points

Little's Law: $W = L / \lambda$

latency = transactions / throughput



Homer wants to minimize response time.

Coz: a Causal Profiler for Linux

(ships with Debian / Ubuntu)



Coz: a Causal Profiler for Linux

(ships with Debian / Ubuntu)

```
> sudo apt install coz-profiler
```

Coz: a Causal Profiler for Linux

(ships with Debian / Ubuntu)

```
> sudo apt install coz-profiler  
> coz run --- ./some_program args
```

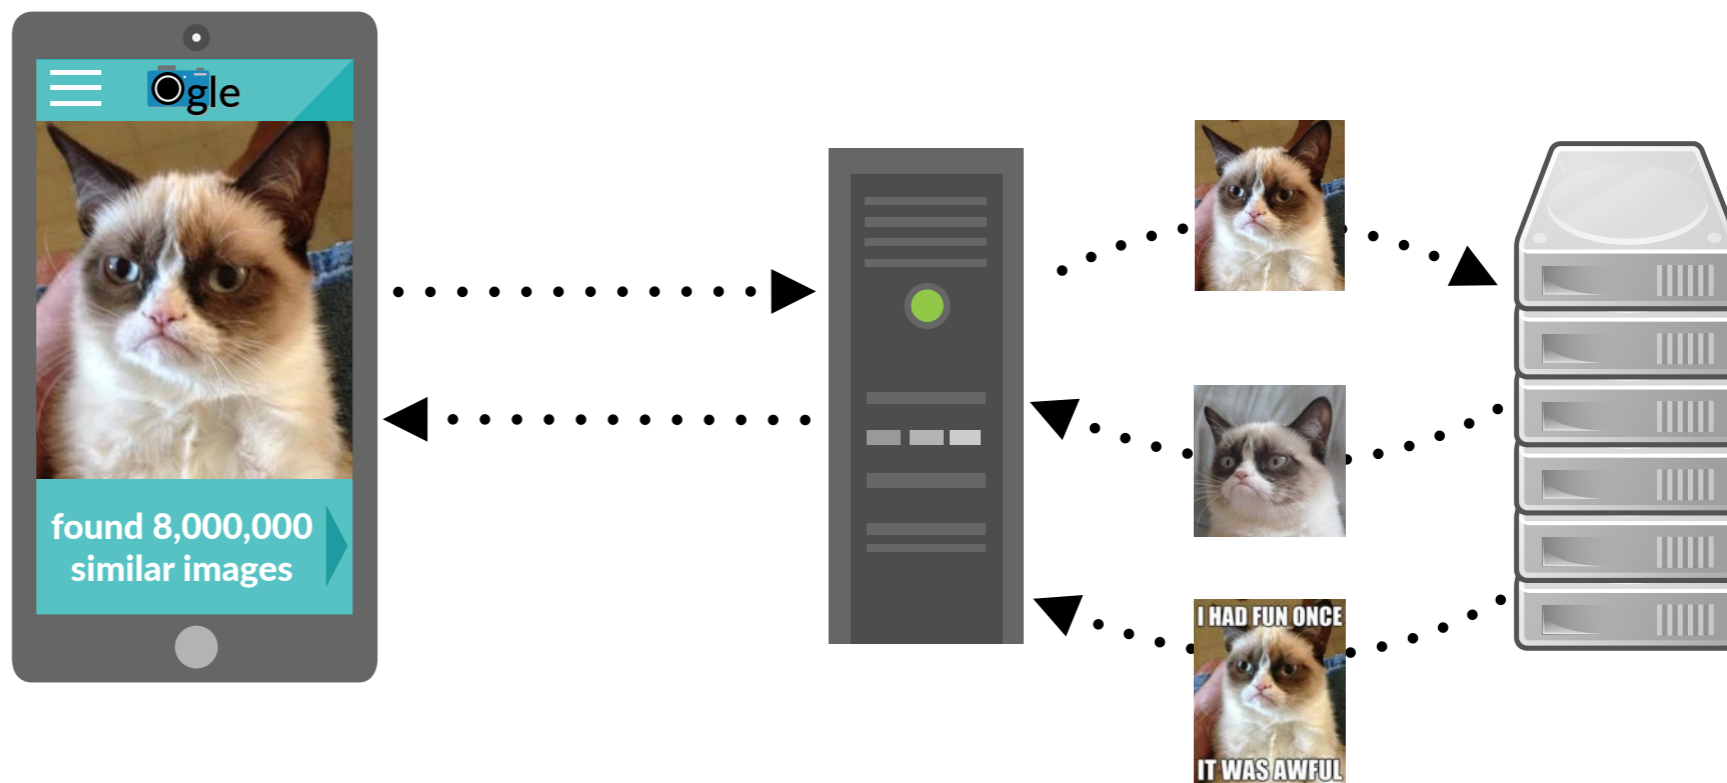

Coz: a Causal Profiler for Linux

(ships with Debian / Ubuntu)


```
> sudo apt install coz-profiler  
> coz run --- ./some_program args
```

Random *performance experiments*


Using Causal Profiling on Ogle



Using Causal Profiling on Ogle




Using Causal Profiling on Ogle




 **PARSEC**
['pär-"sek] A unit of measure

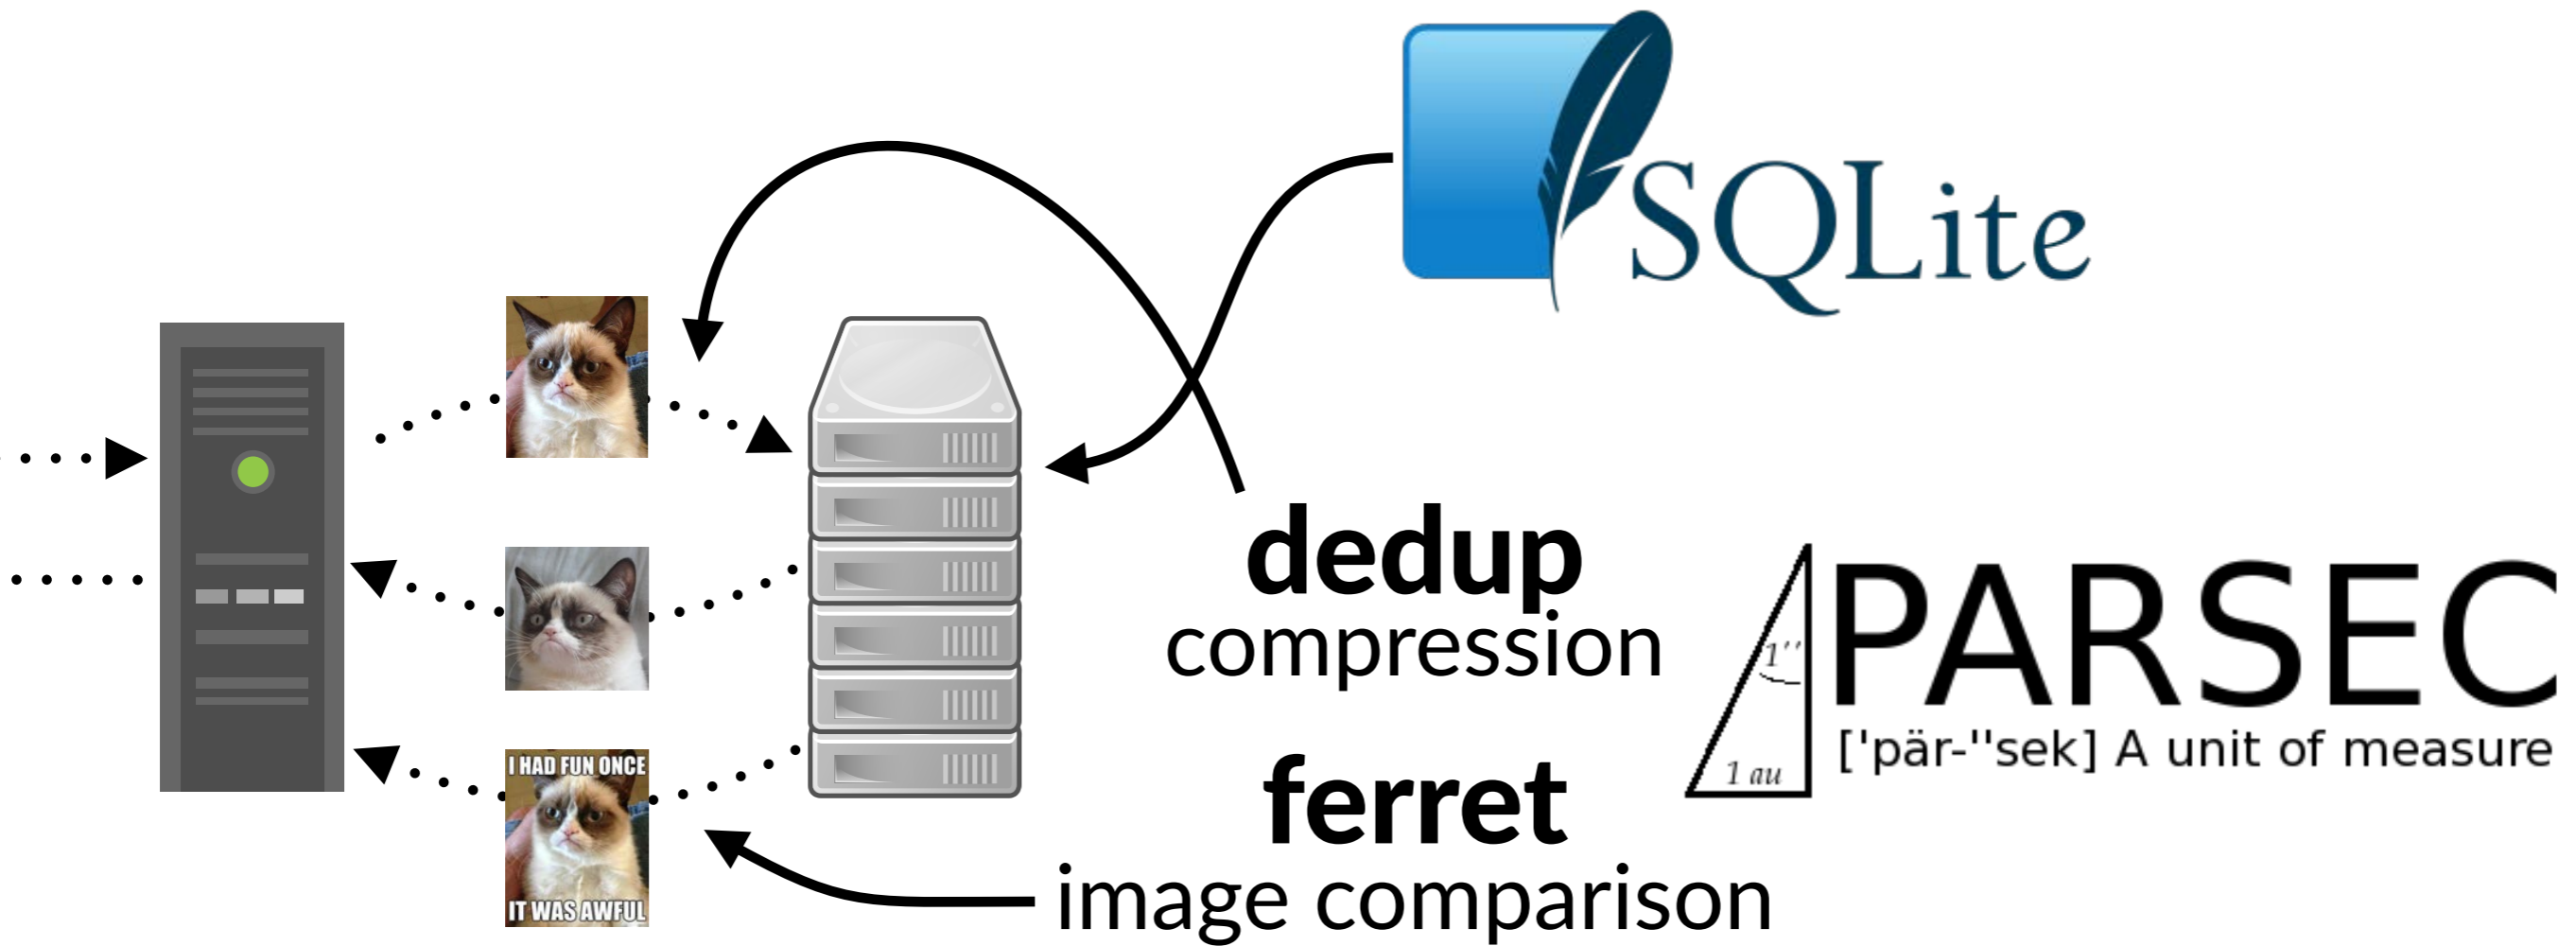
Using Causal Profiling on Ogle



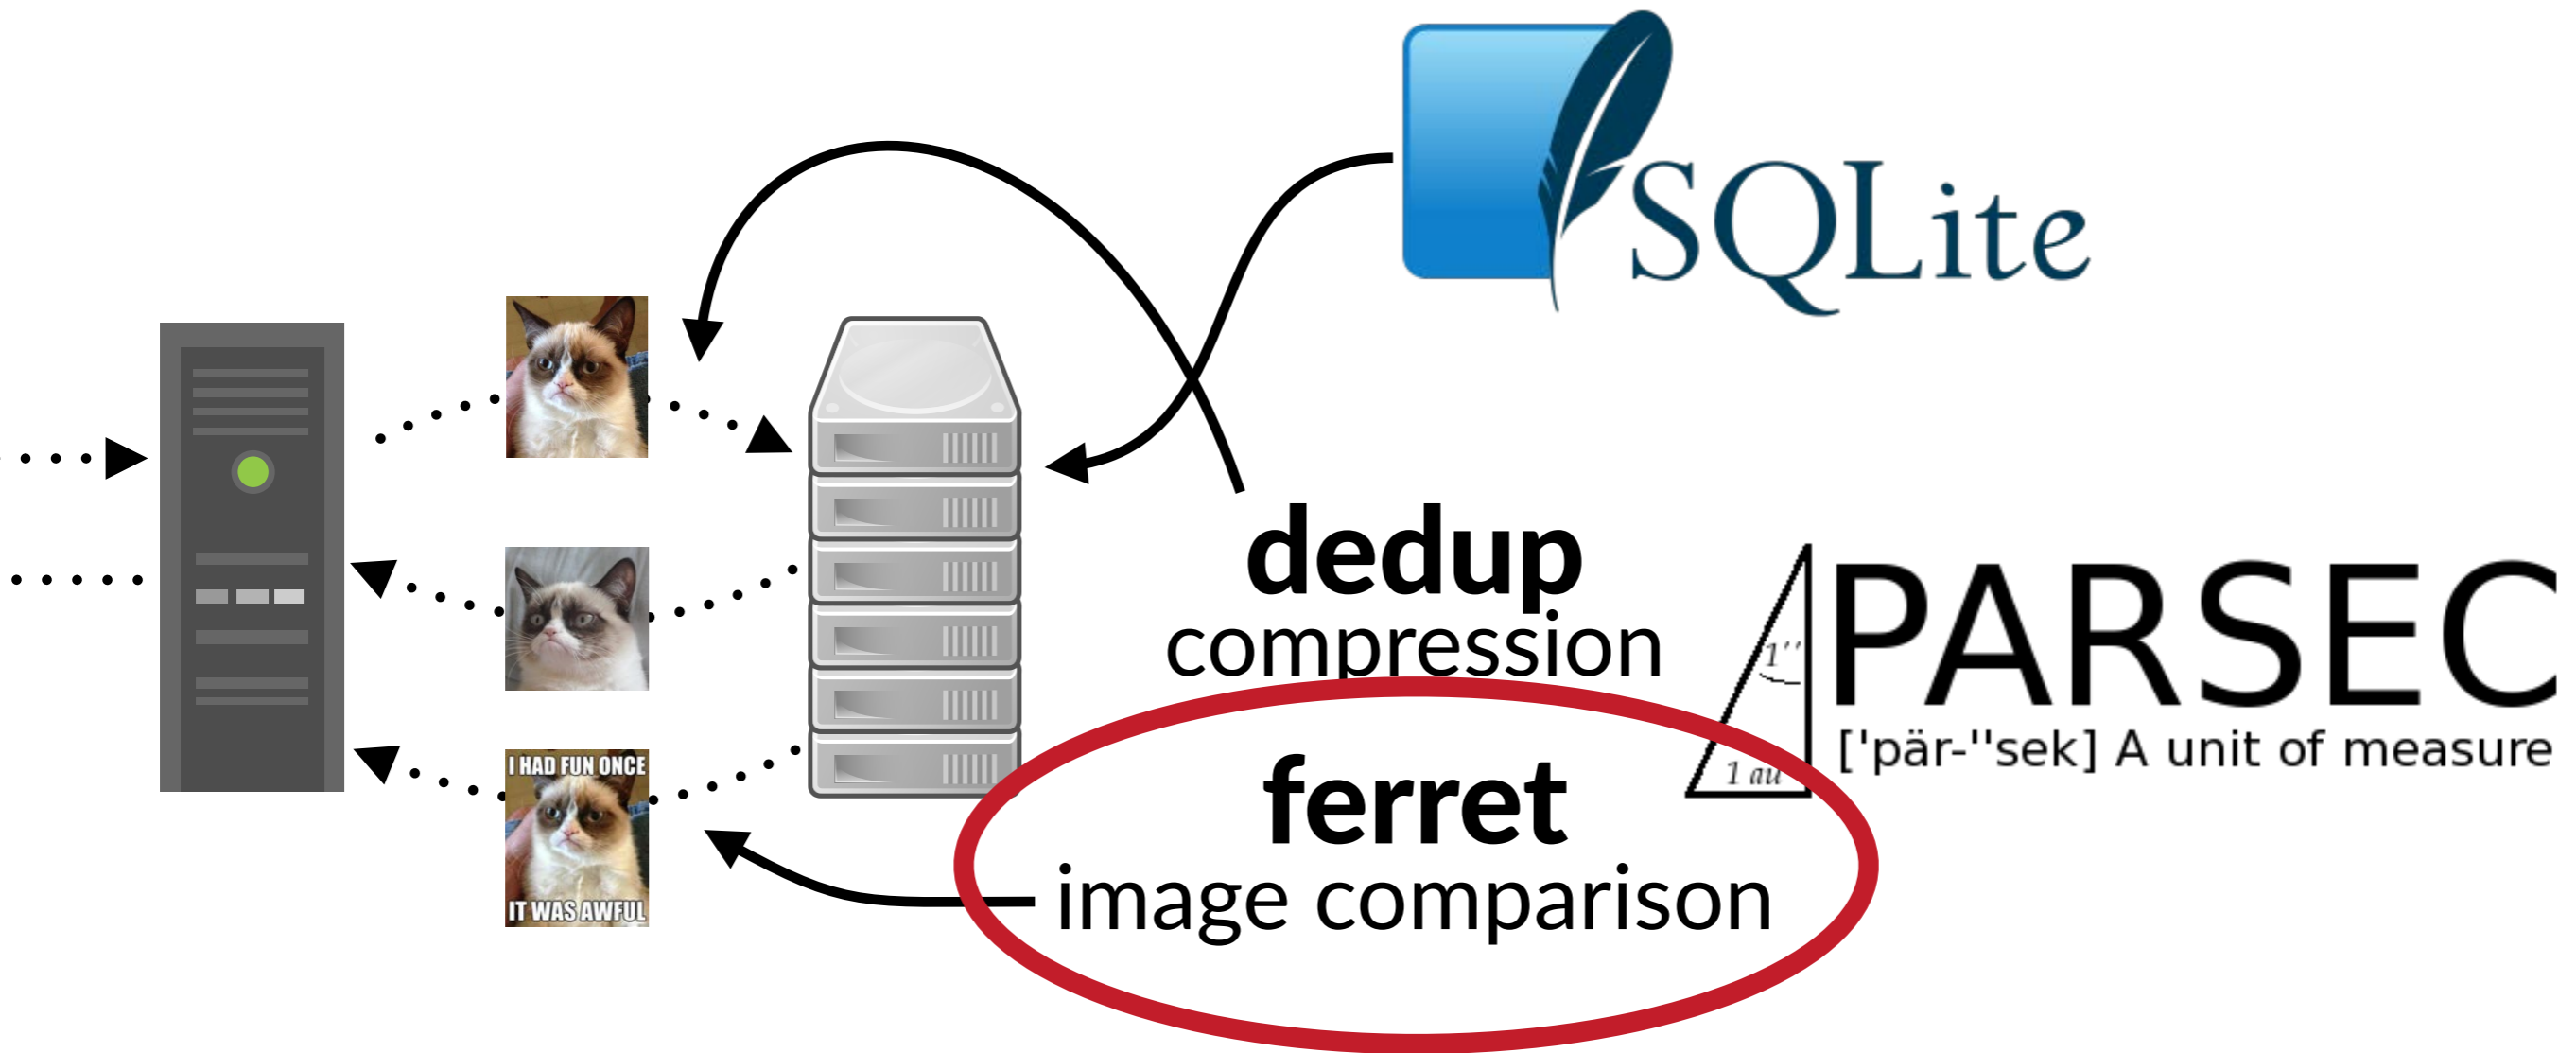
Using Causal Profiling on Ogle




Using Causal Profiling on Ogle



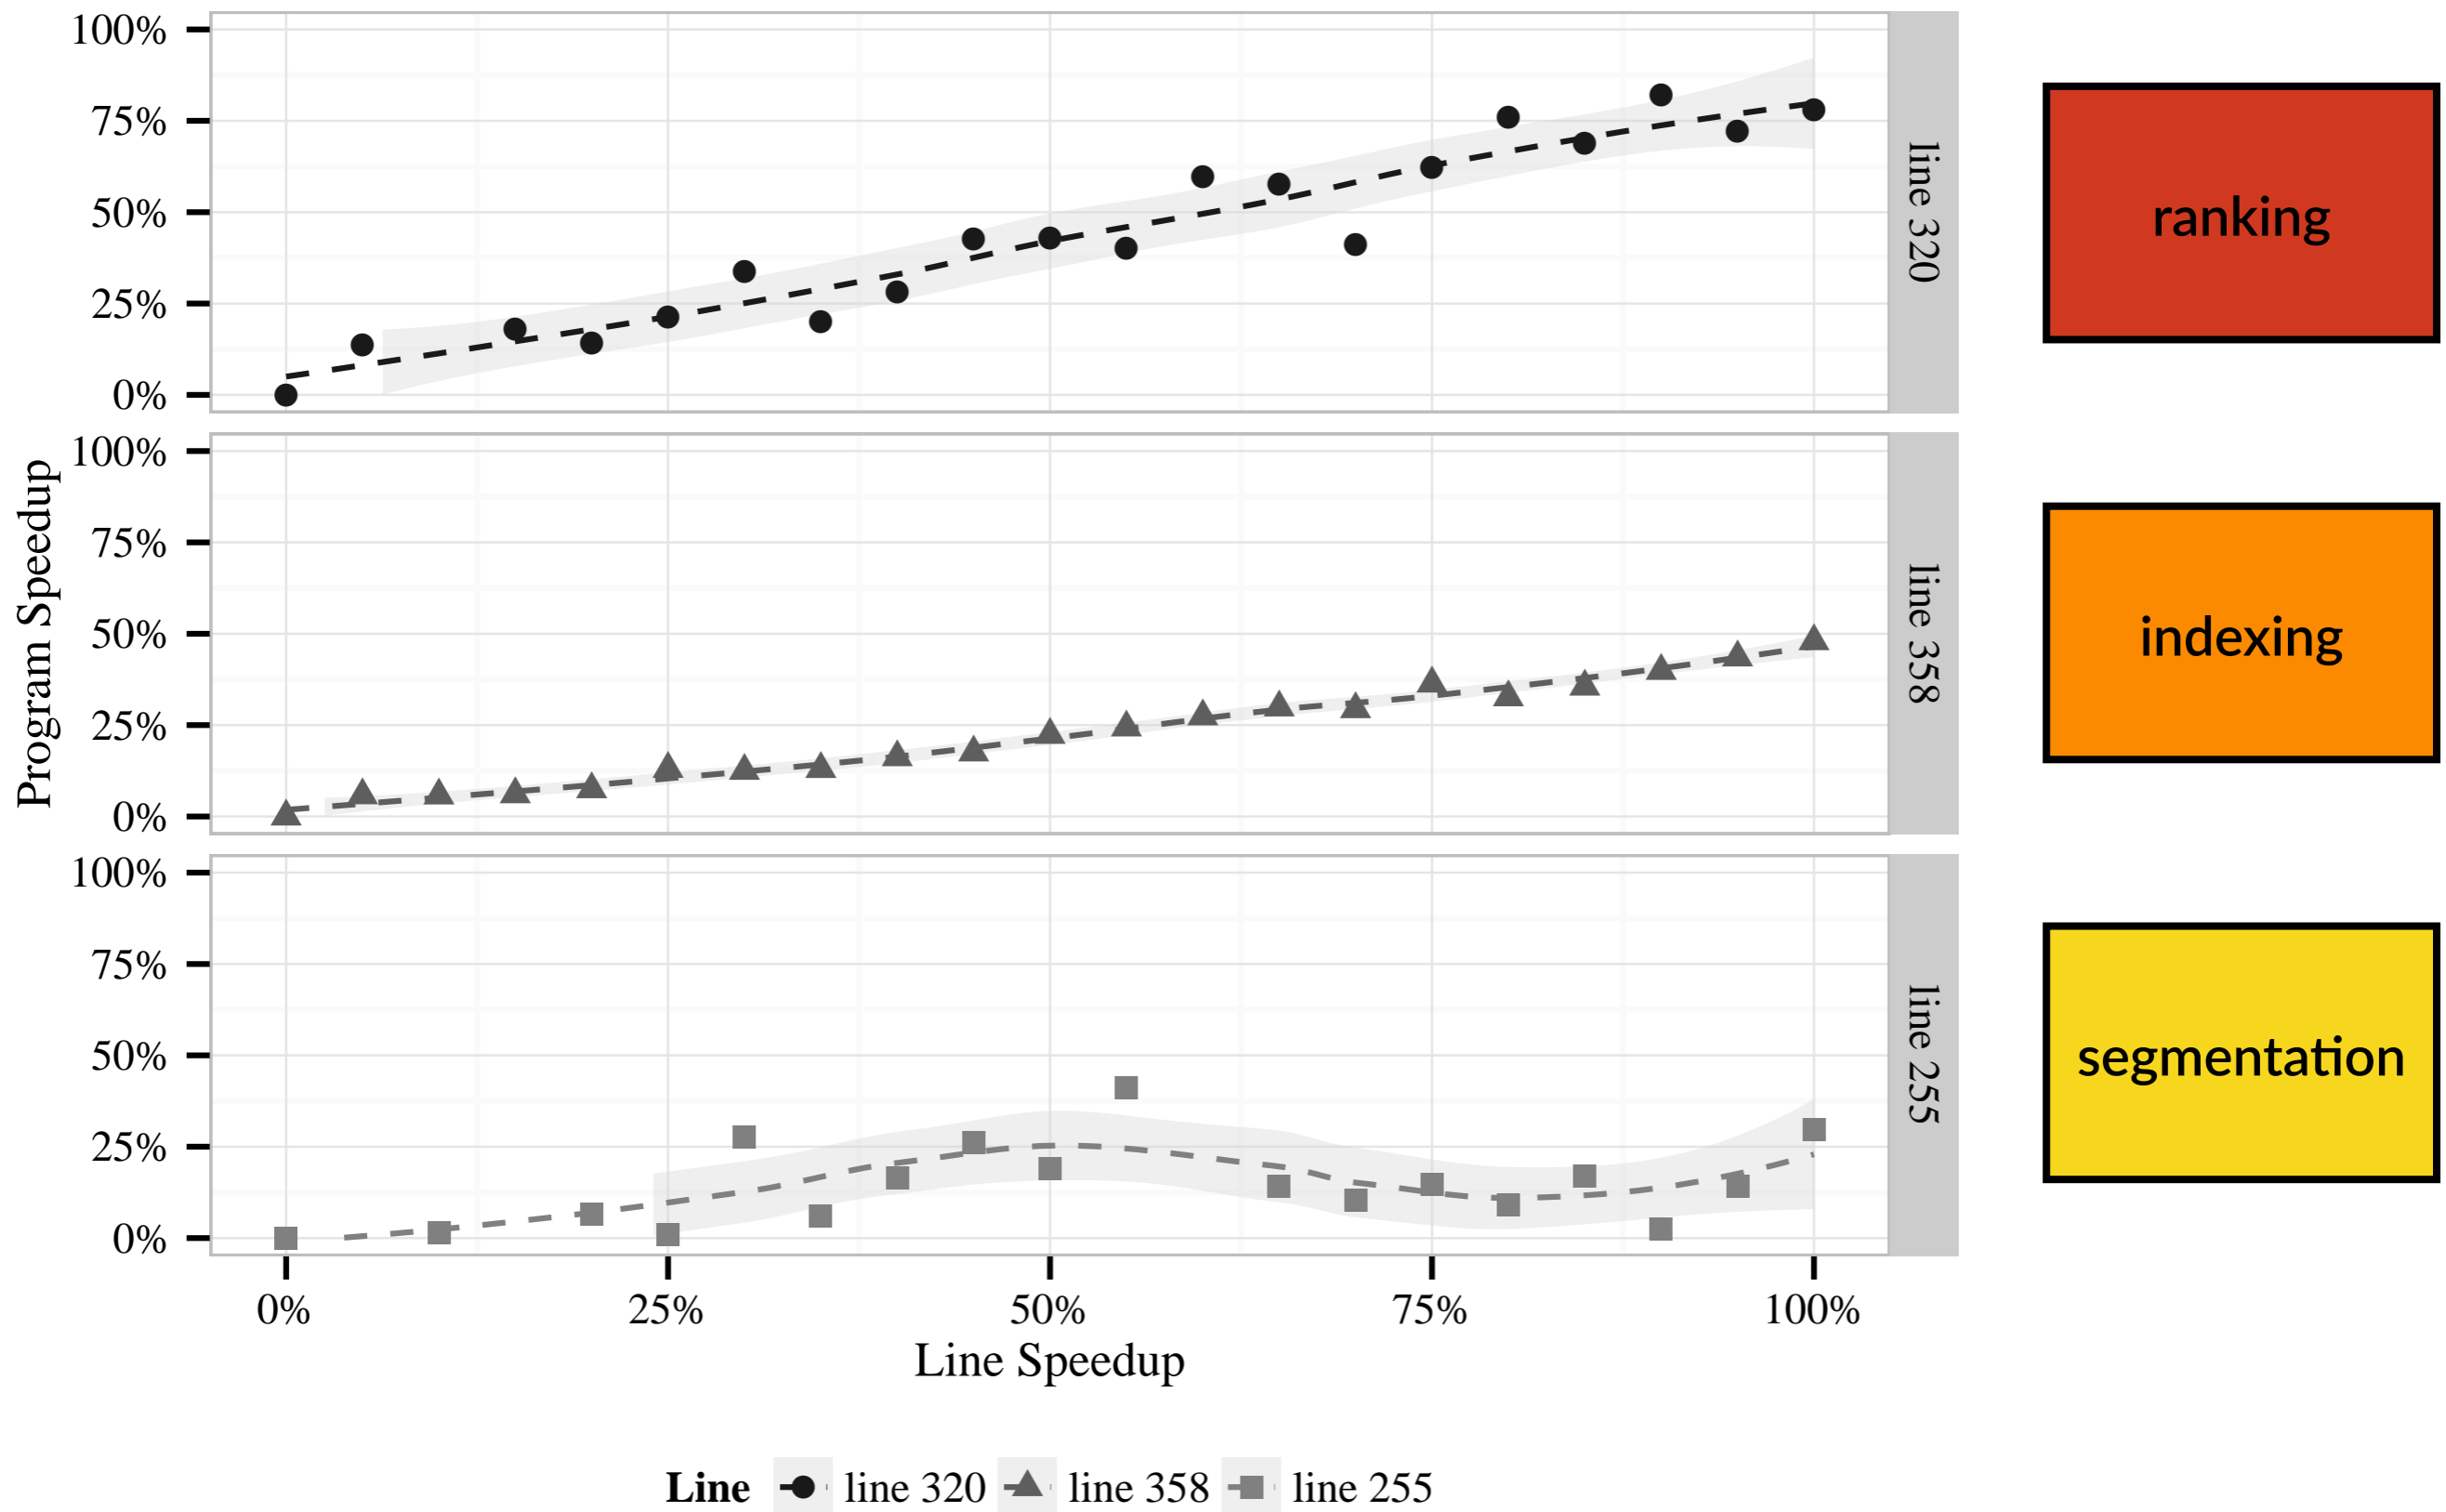
Using Causal Profiling on Ogle



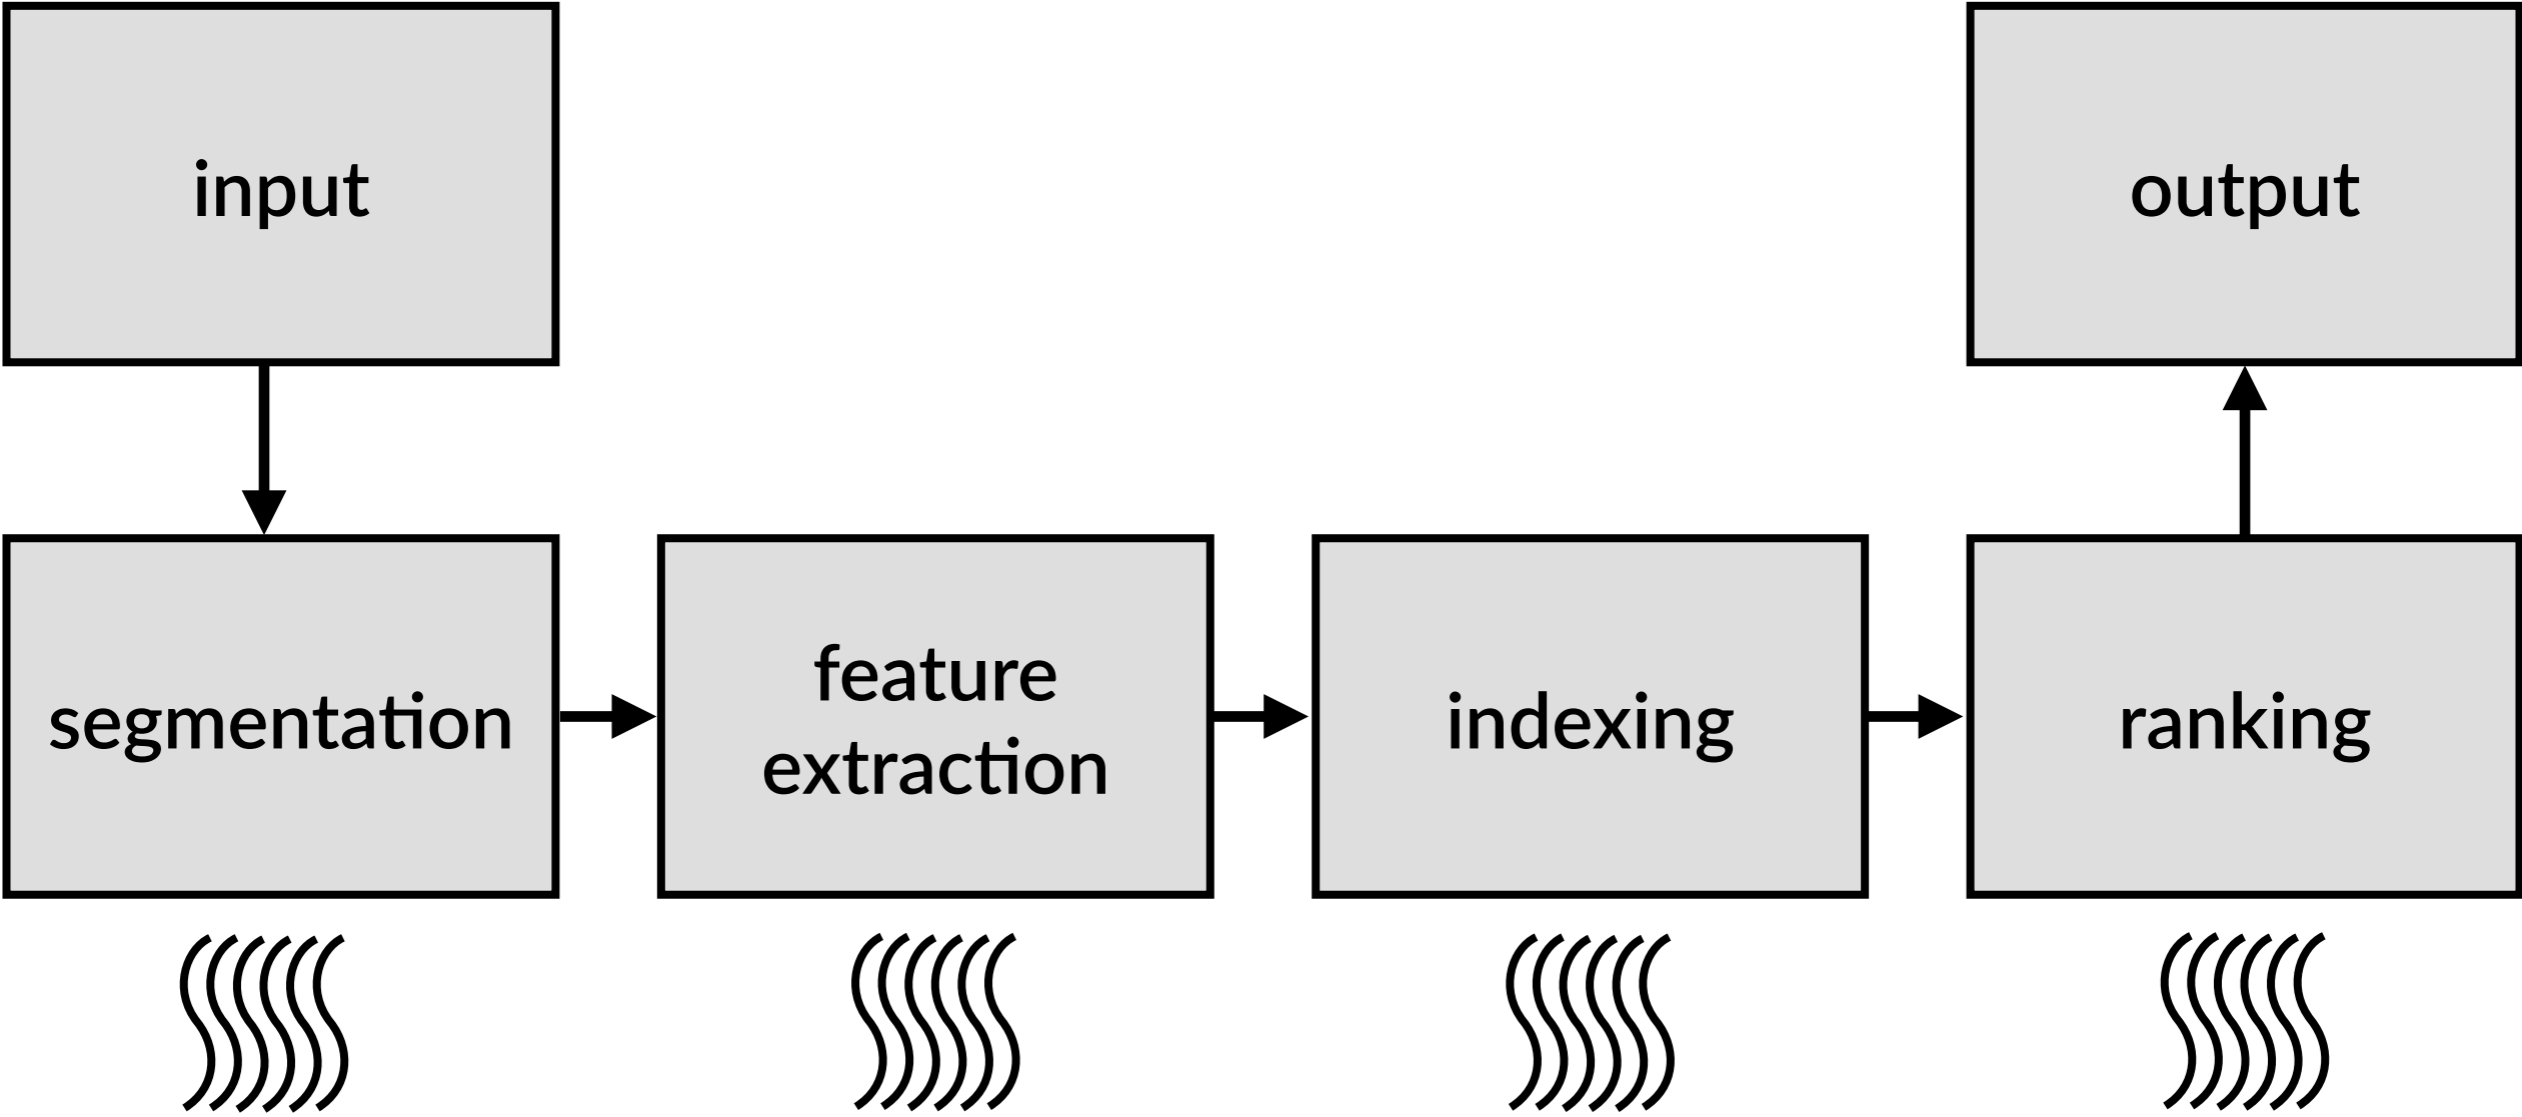
Ferret



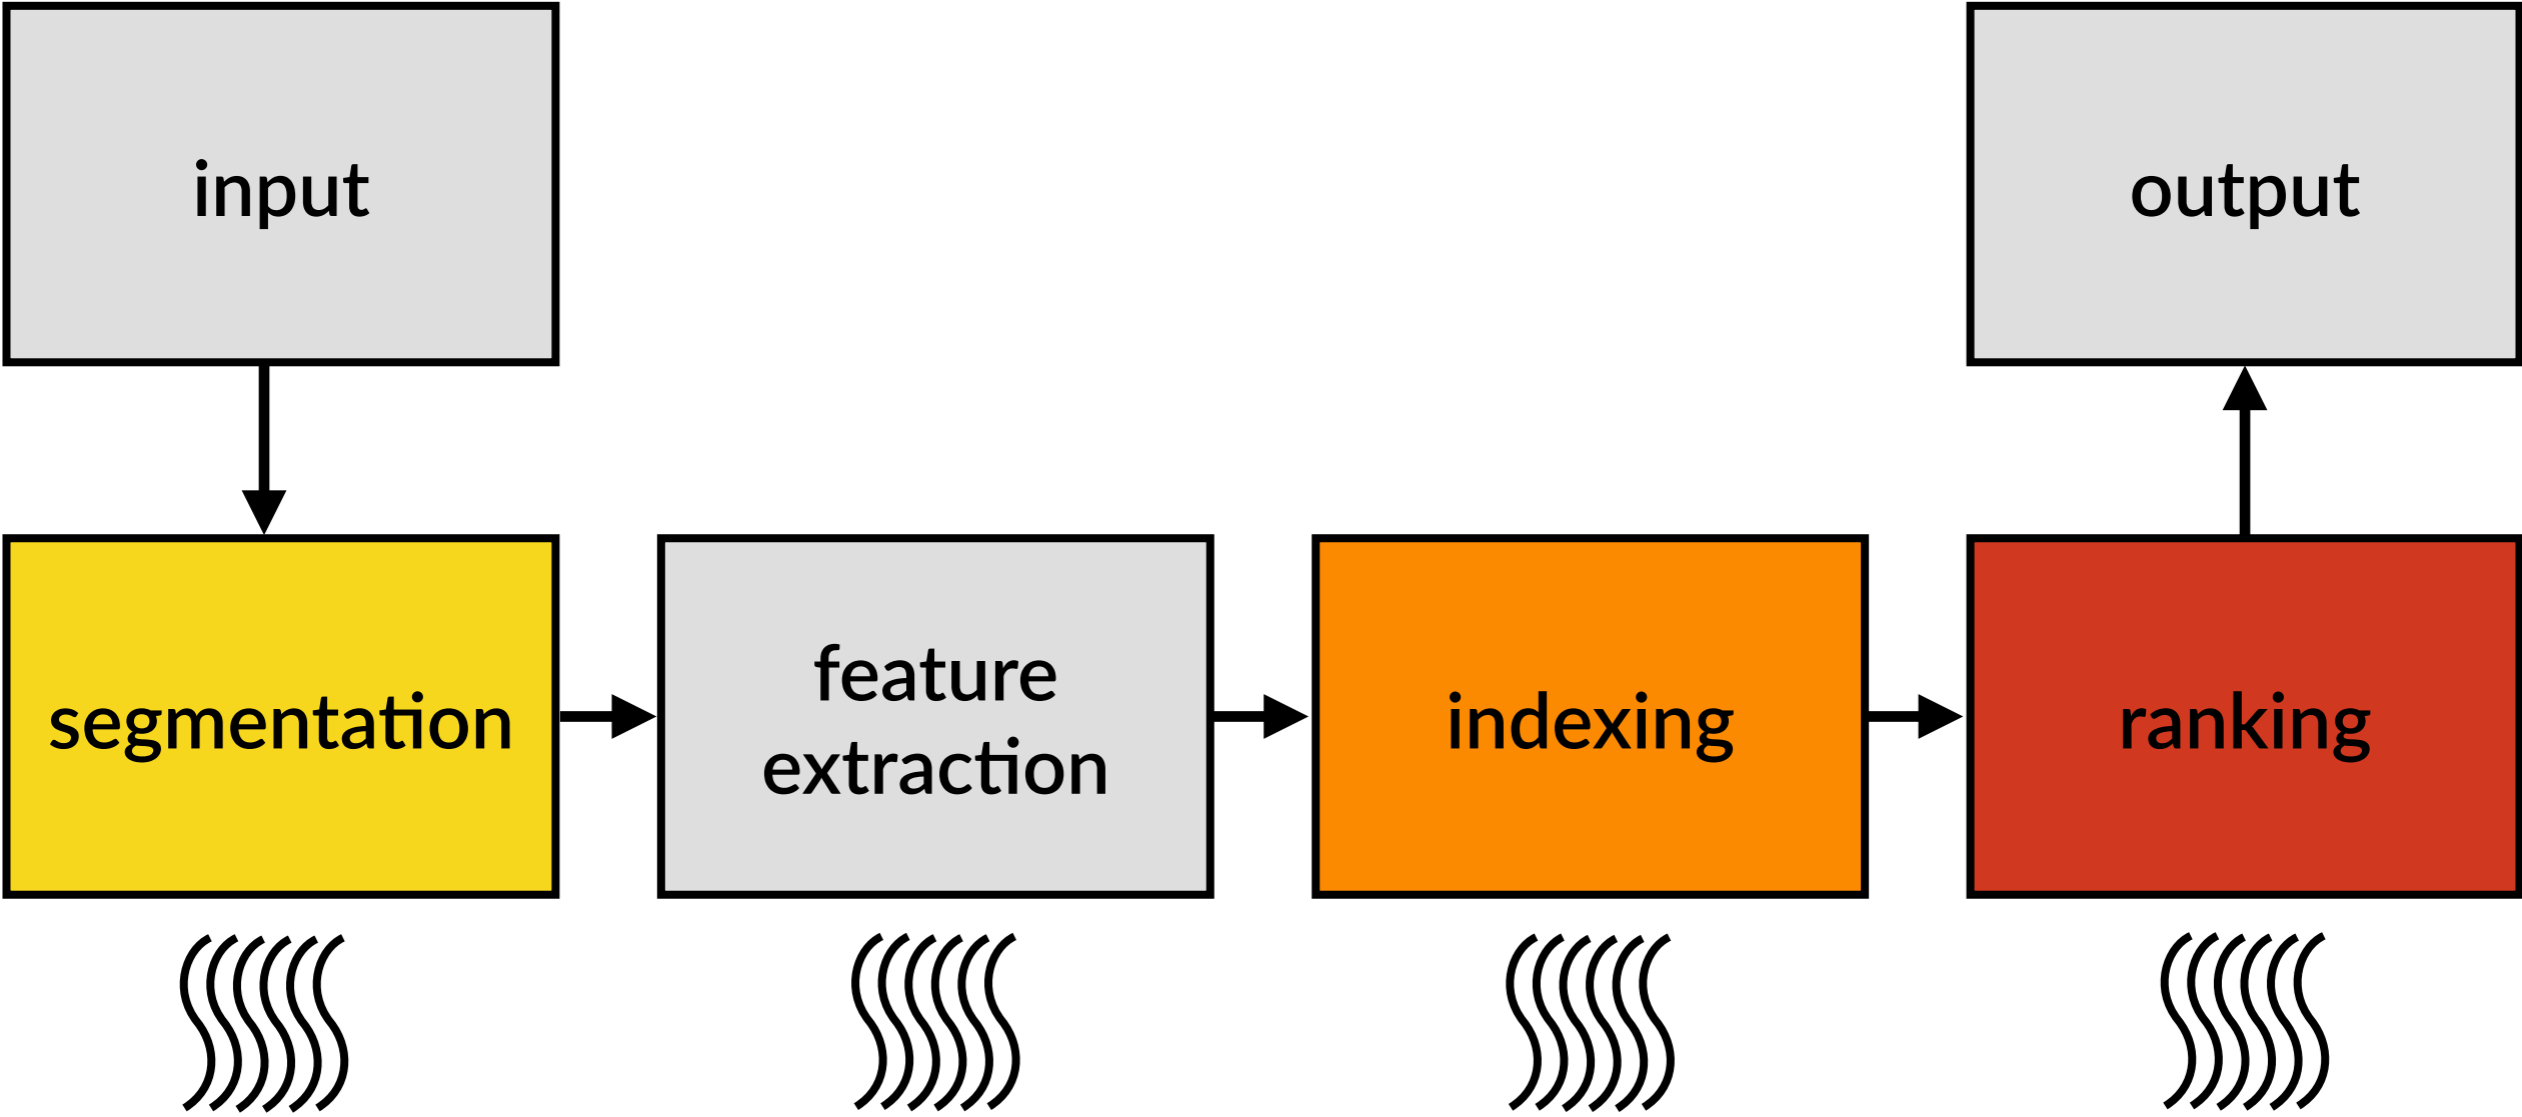
Ferret



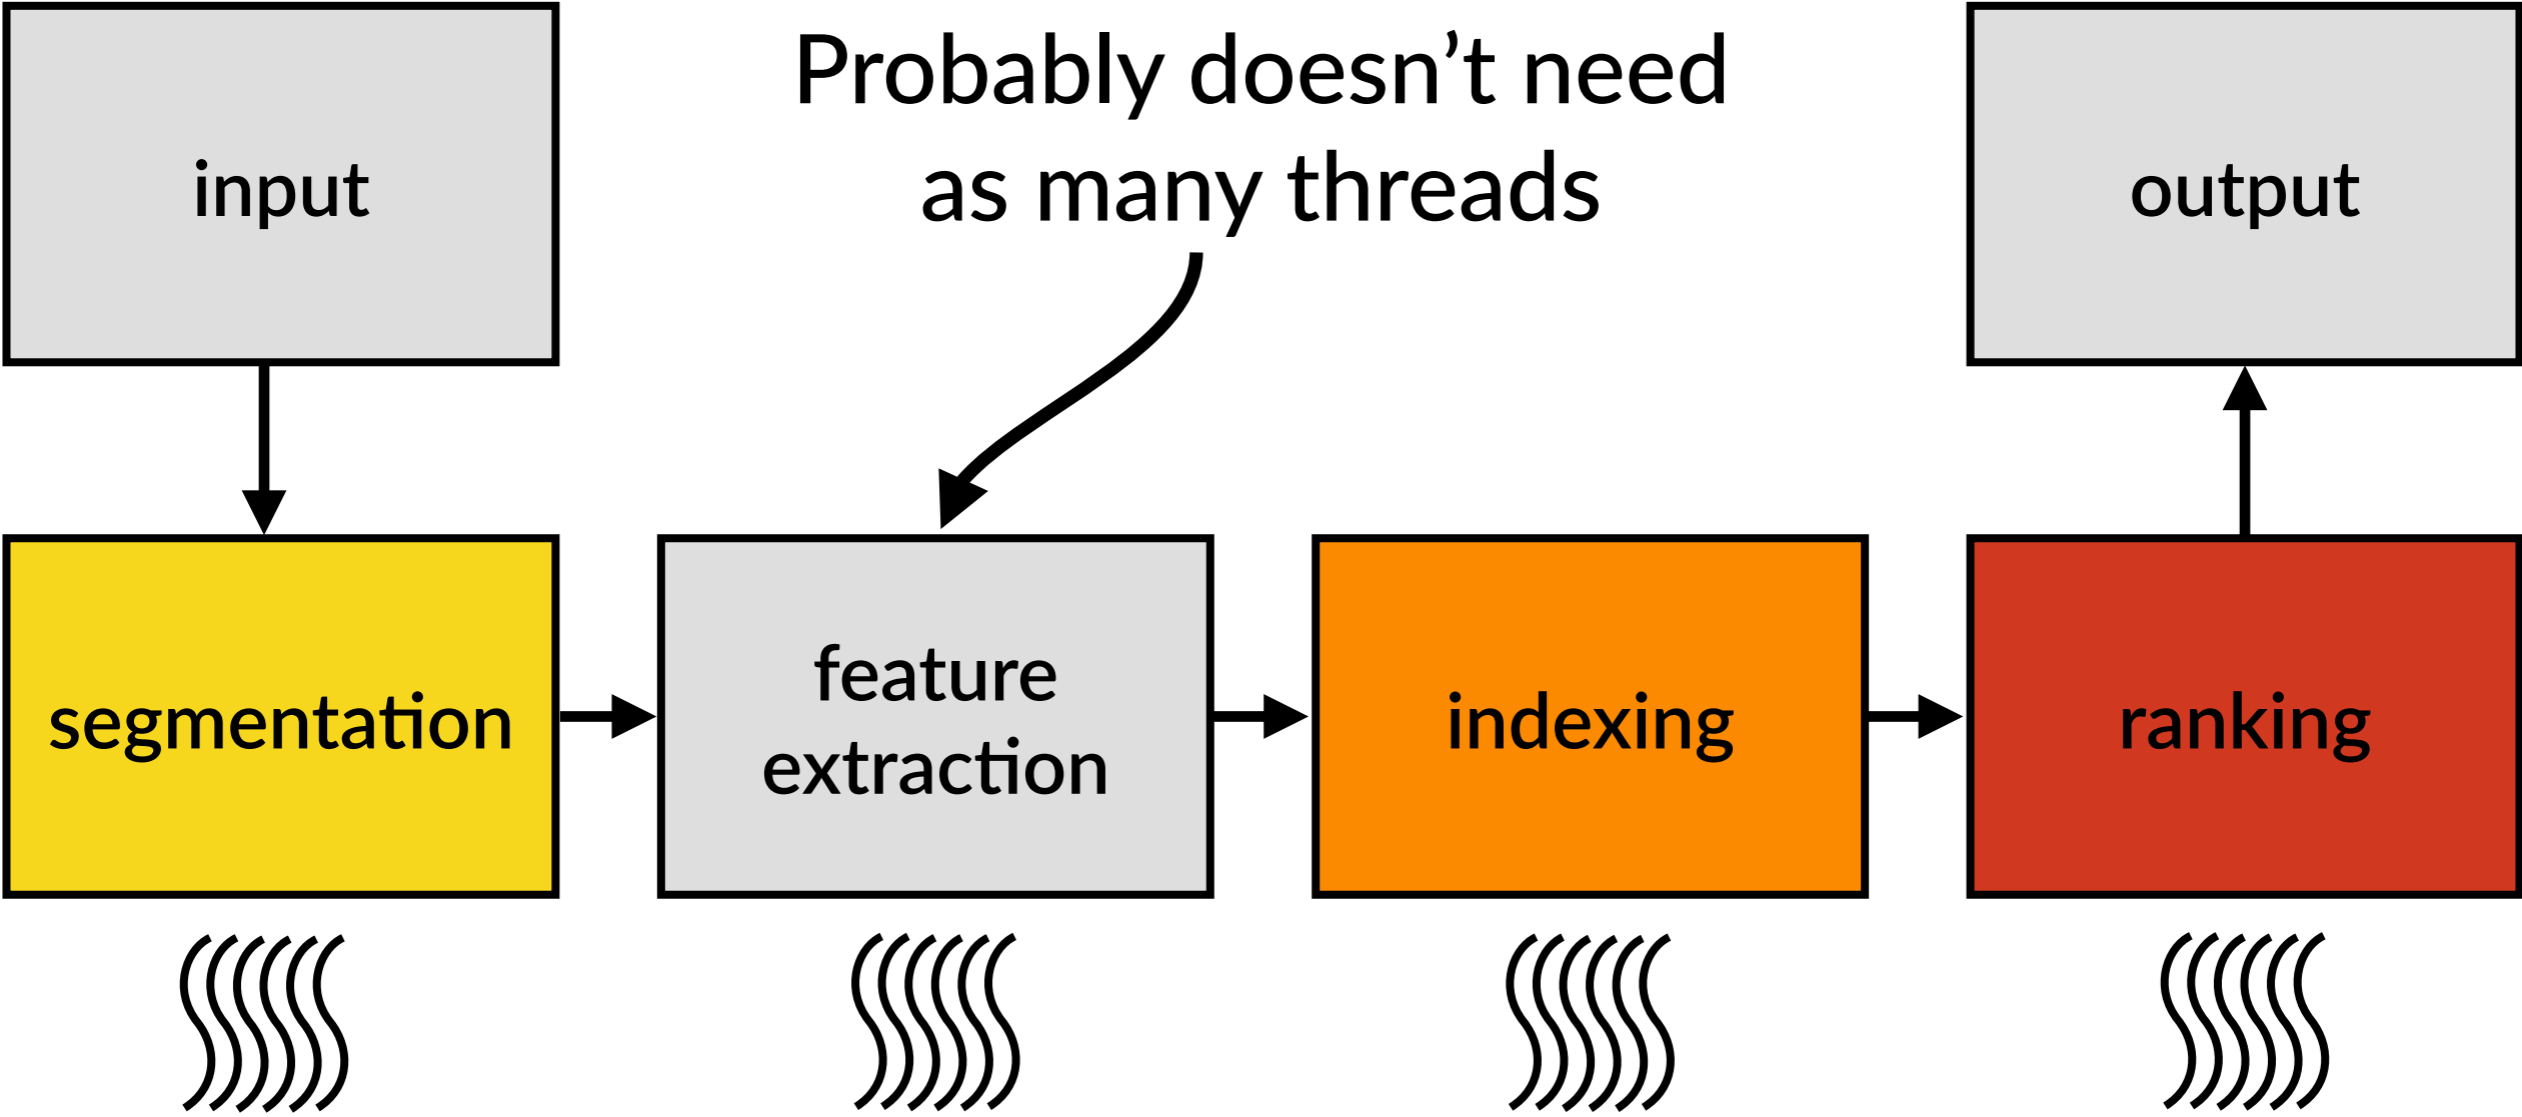
Ferret



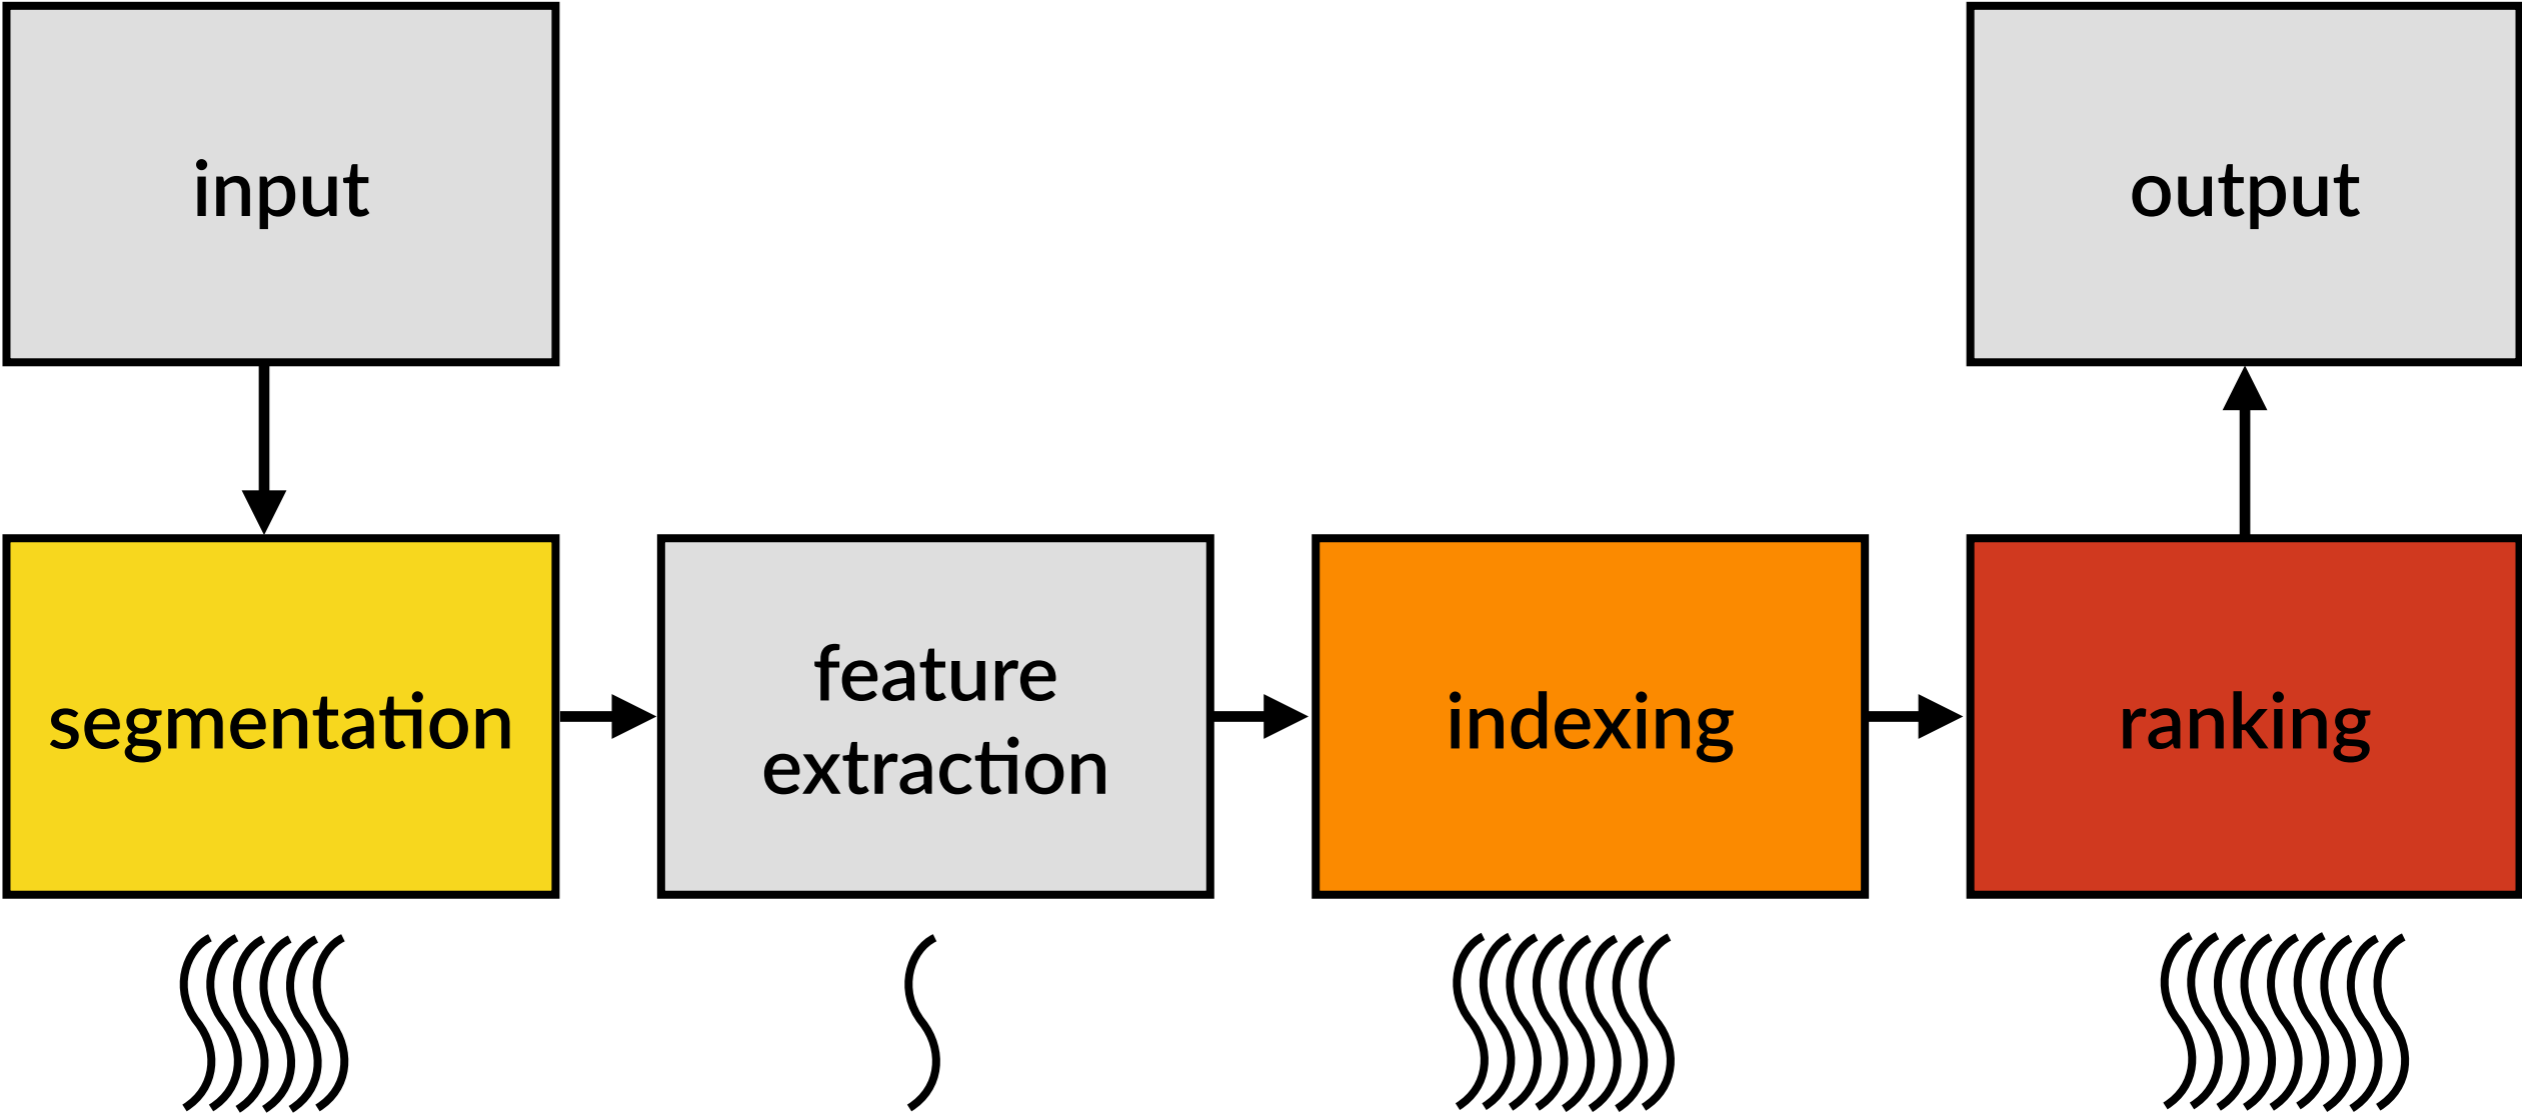
Ferret



Ferret




Ferret



Ferret

21% Speedup



What did Coz predict?



ranking

What did Coz predict?



Increased from 16 to 22 threads


What did Coz predict?



Increased from 16 to 22 threads

27% increase in ranking throughput


What did Coz predict?



ranking


27% increase in ranking throughput

What did Coz predict?




27% increase in ranking throughput

What did Coz predict?



27% increase in ranking throughput


What did Coz predict?



ranking

27% increase in ranking throughput
Coz predicted a 21% improvement

What did Coz predict?

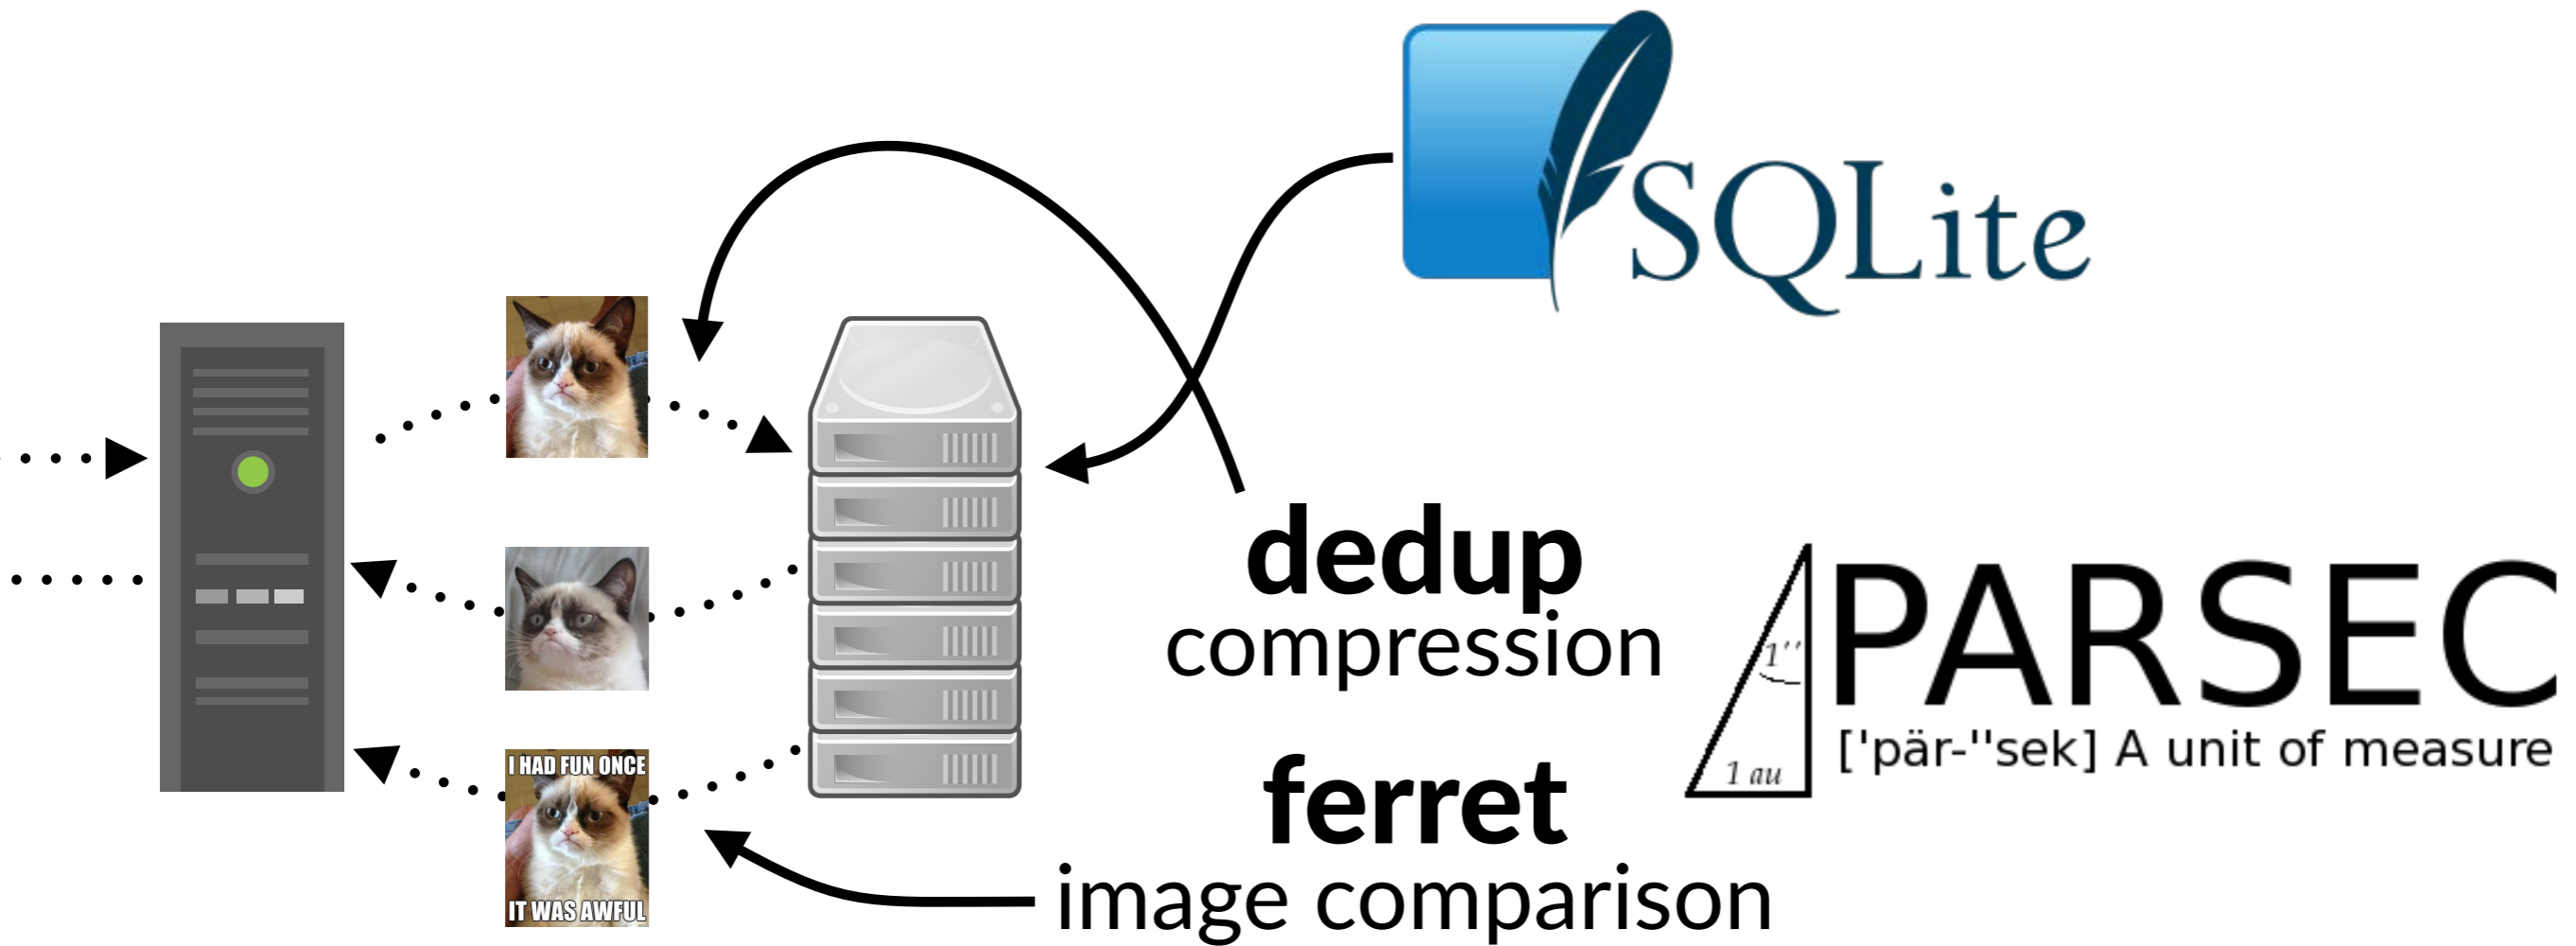


27% increase in ranking throughput

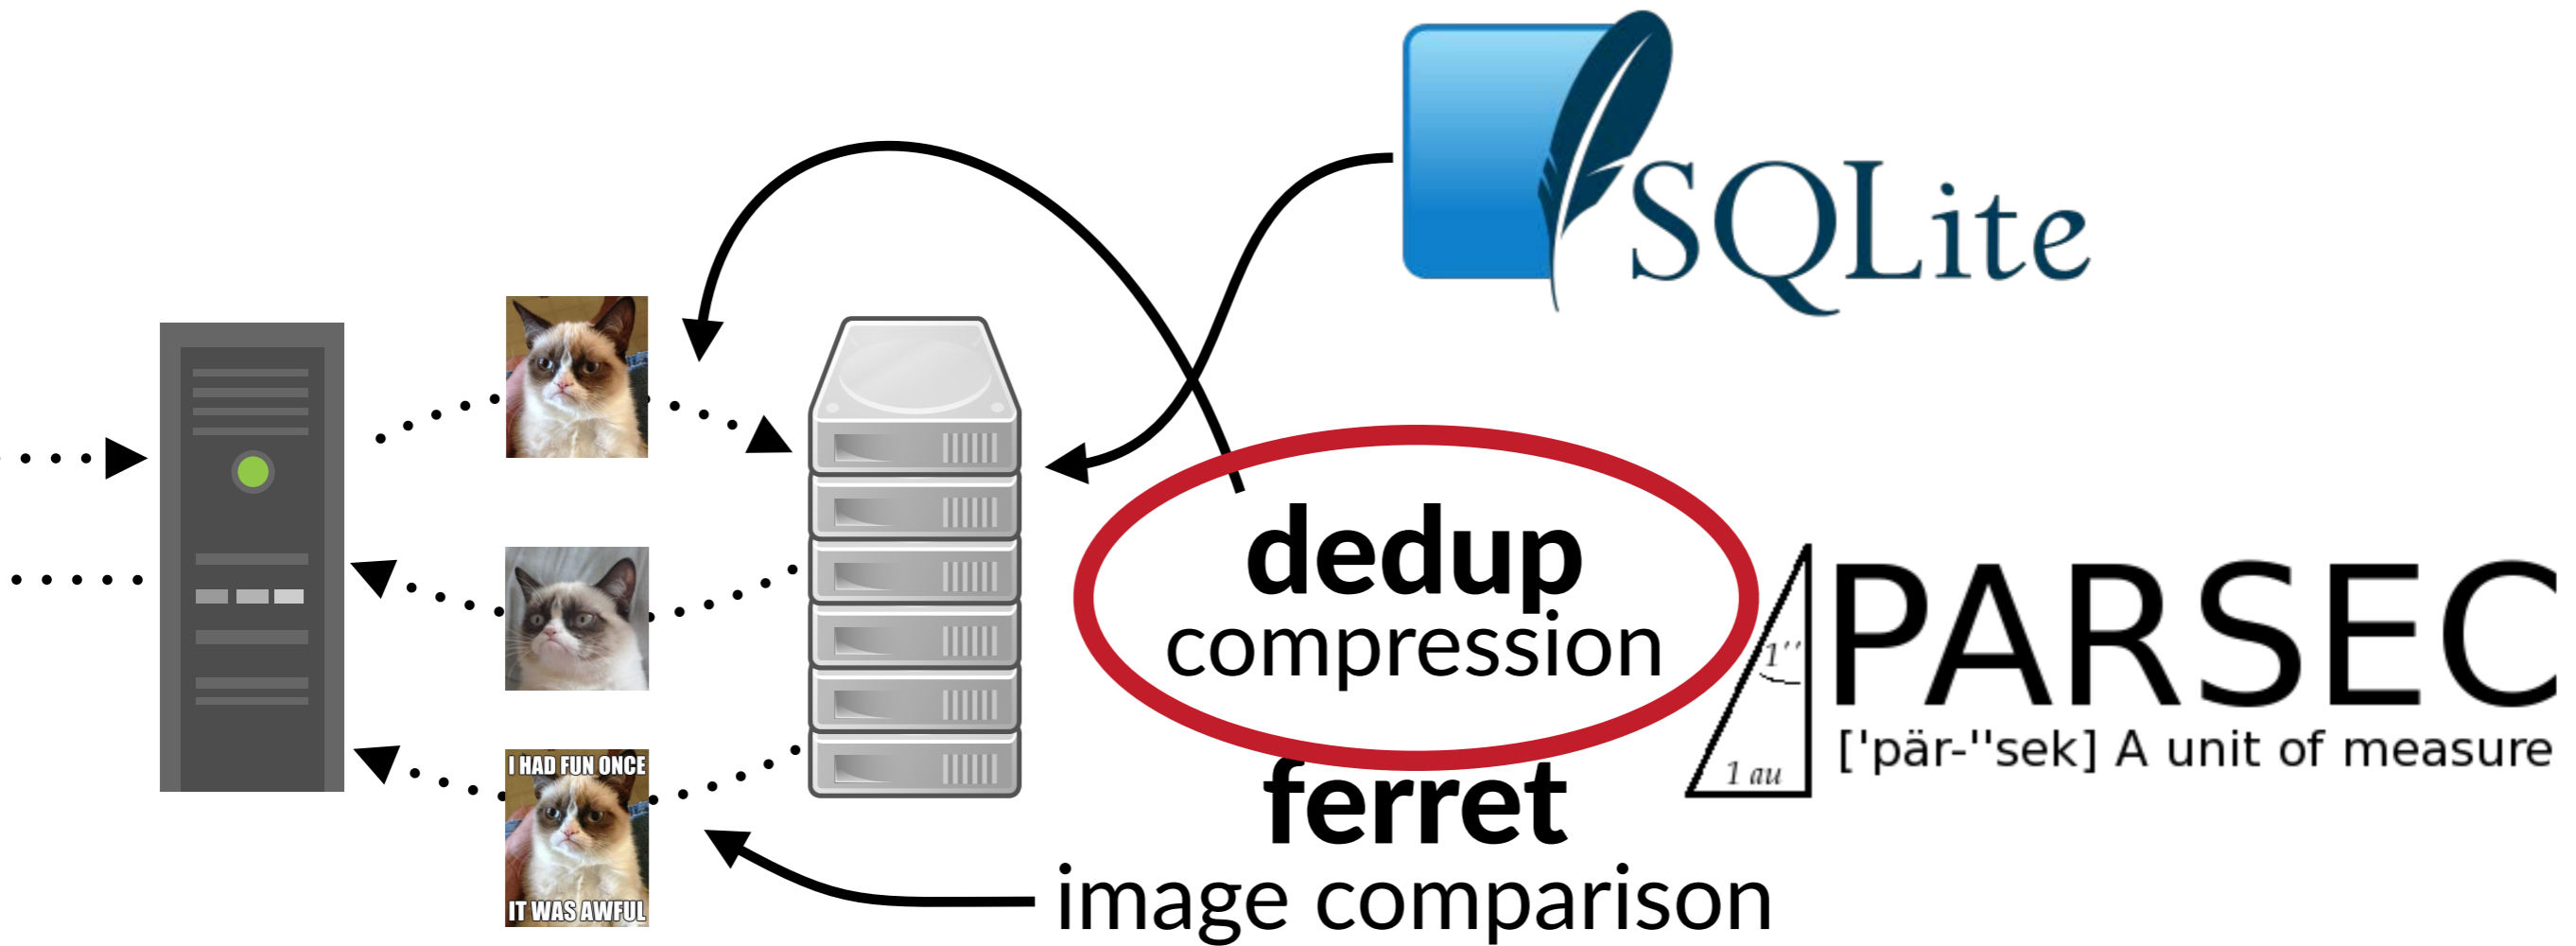
Coz predicted a 21% improvement

Exactly what we observed

Using Causal Profiling on Ogle



Using Causal Profiling on Ogle




Dedup

Compression via deduplication



Dedup

Compression via deduplication




Dedup

Compression via deduplication






Dedup

Compression via deduplication



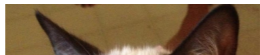

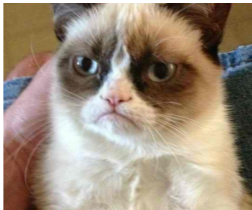
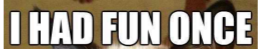

Dedup

Compression via deduplication



Dedup


Compression via deduplication

	 			 	
--	--	---	--	--	--

Dedup

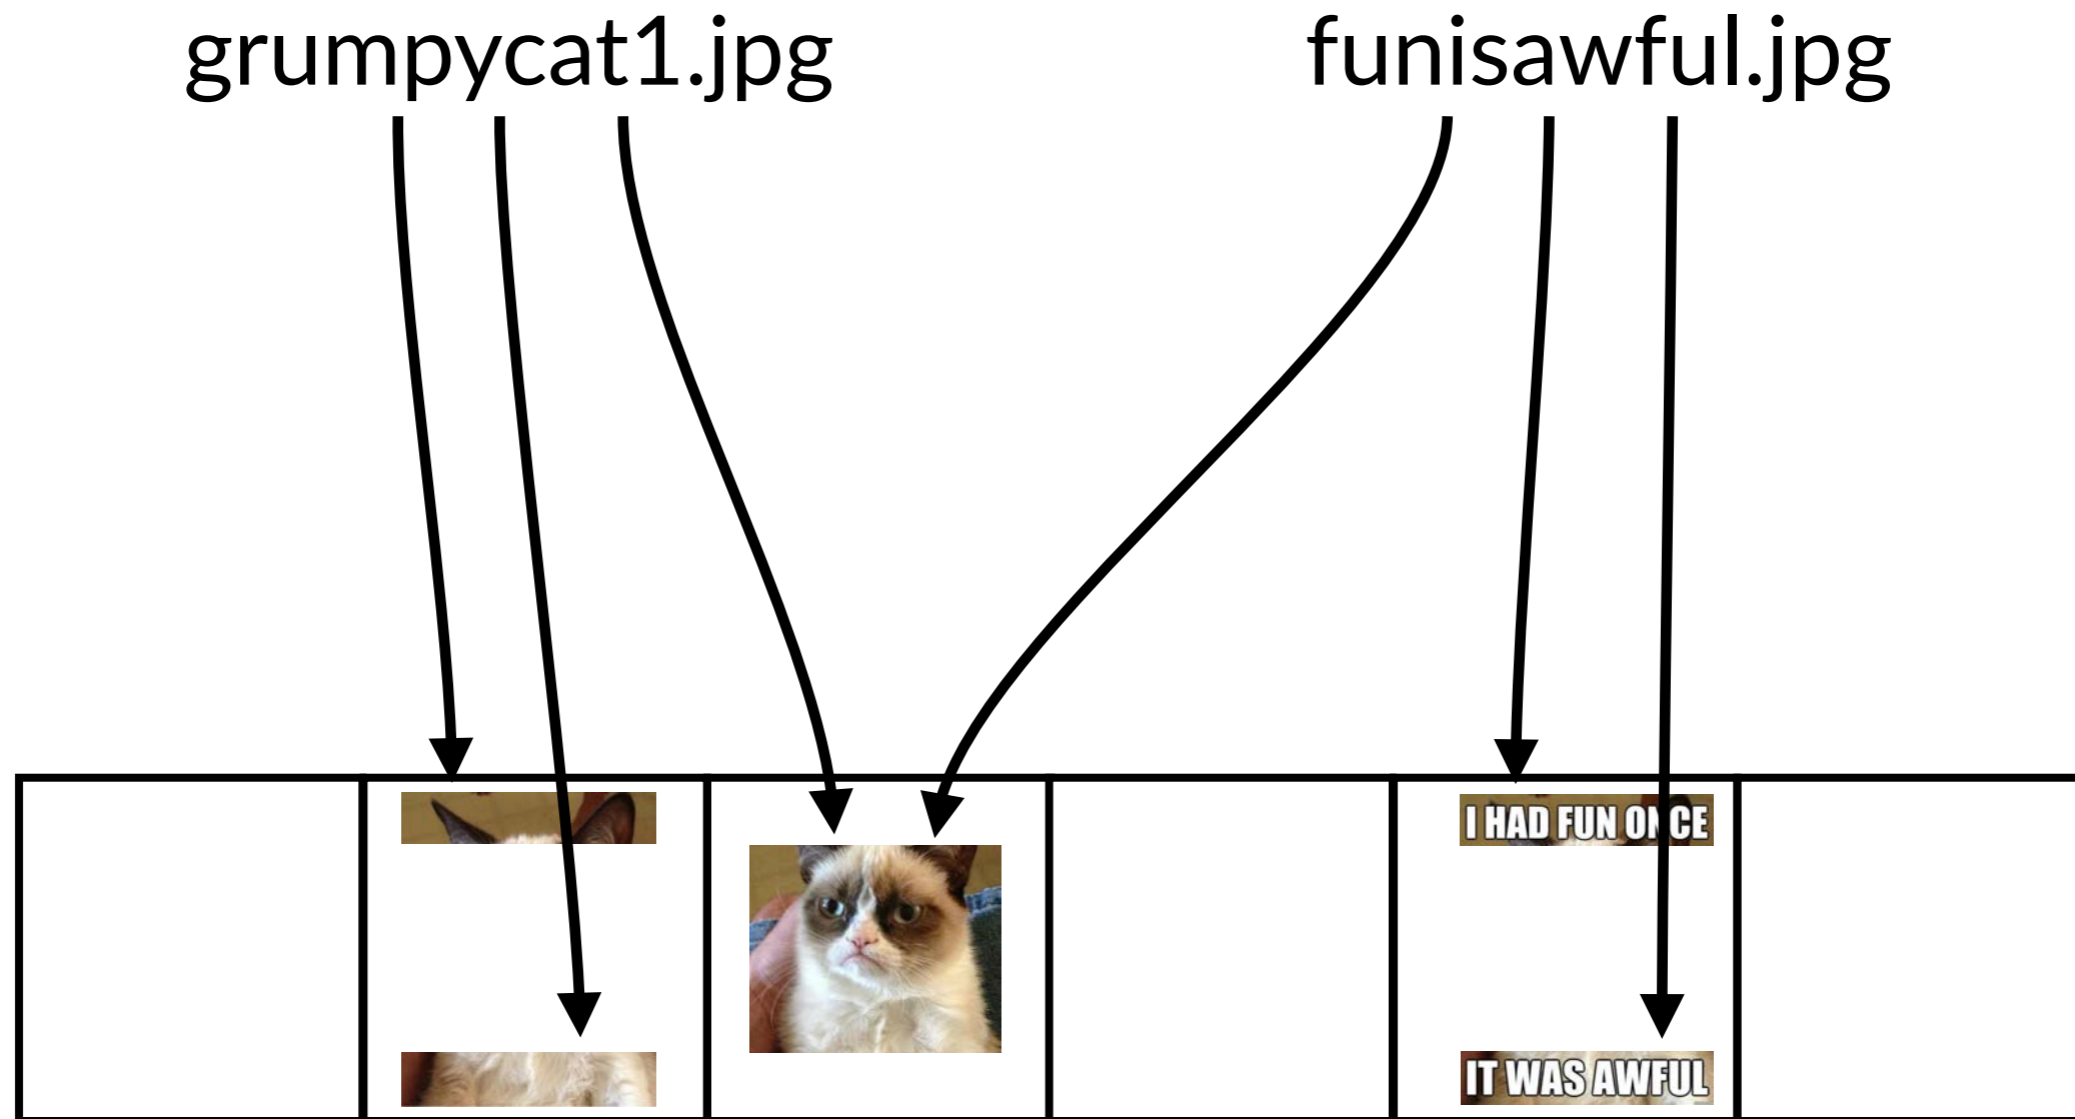
Compression via deduplication

grumpycat1.jpg




Dedup

Compression via deduplication



Dedup


Compression via deduplication



Dedup

Compression via deduplication

hash_function (





)



Dedup

Compression via deduplication



$i = \text{hash_function}(\text{img})$



Dedup

Compression via deduplication

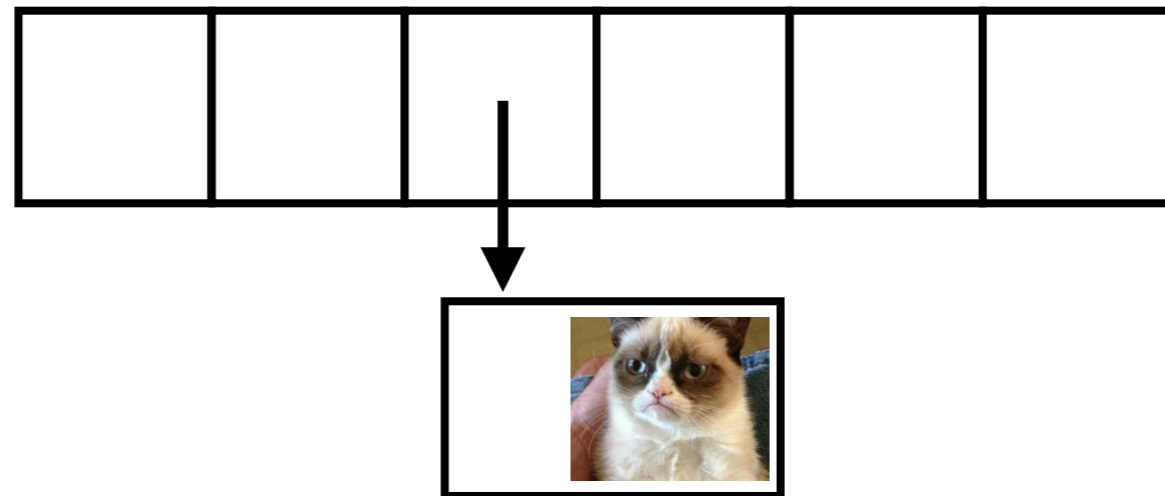
`i = hash_function()`



Dedup

Compression via deduplication




hash_function()



Dedup

Compression via deduplication

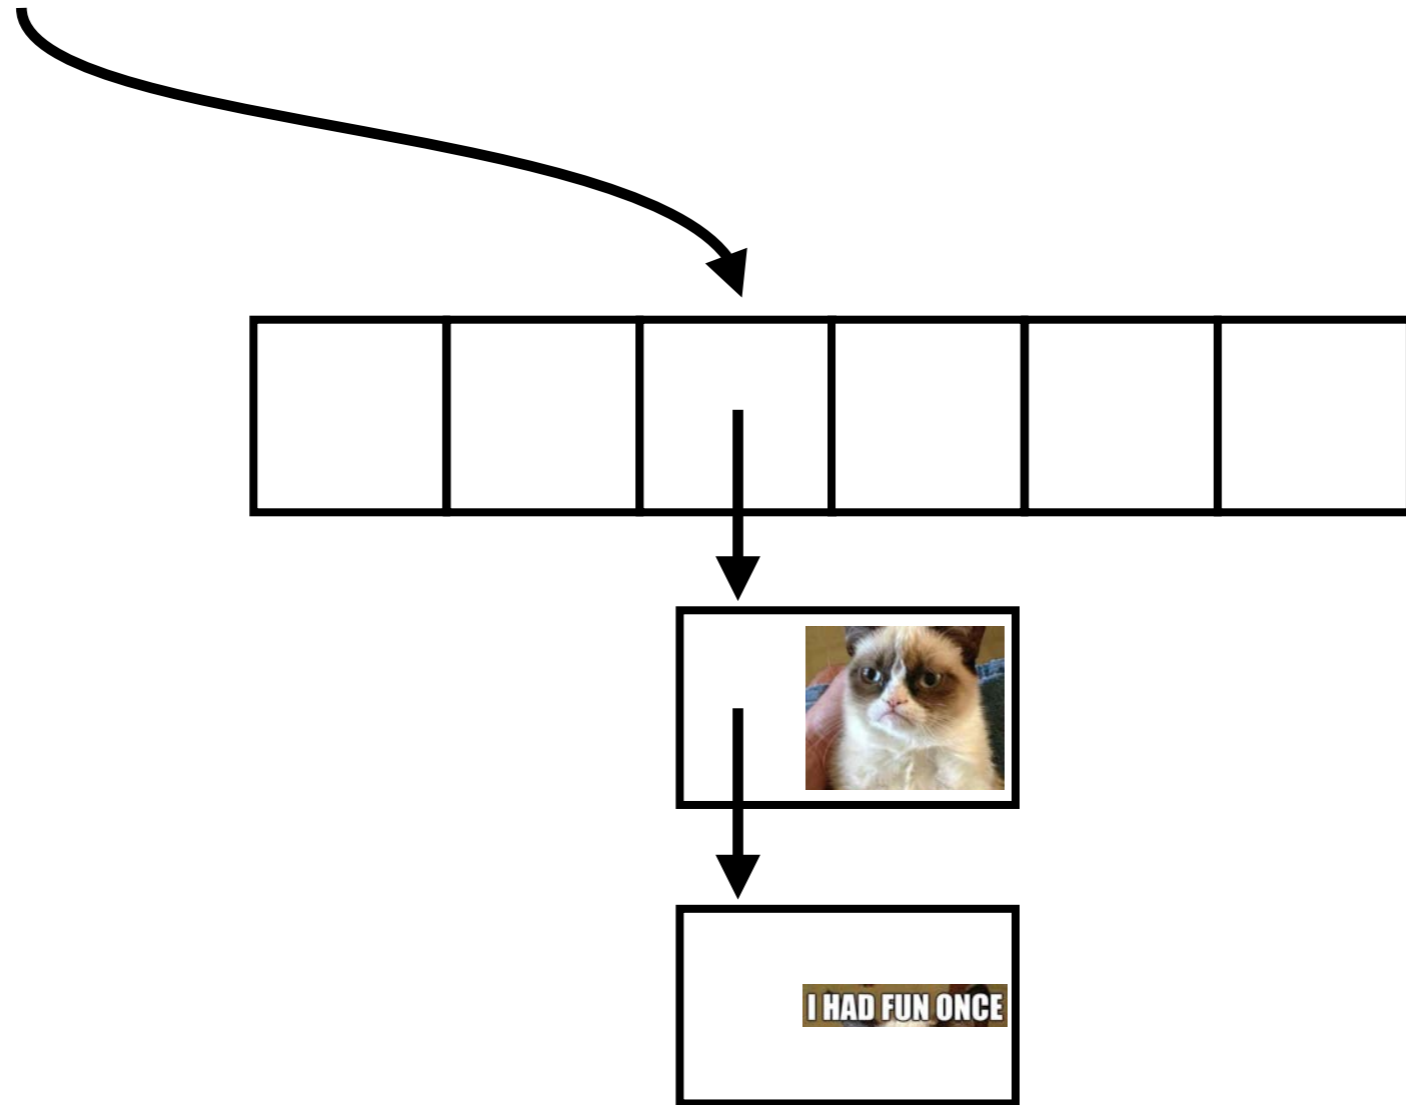
$i = \text{hash_function}(\text{I HAD FUN ONCE})$



Dedup

Compression via deduplication

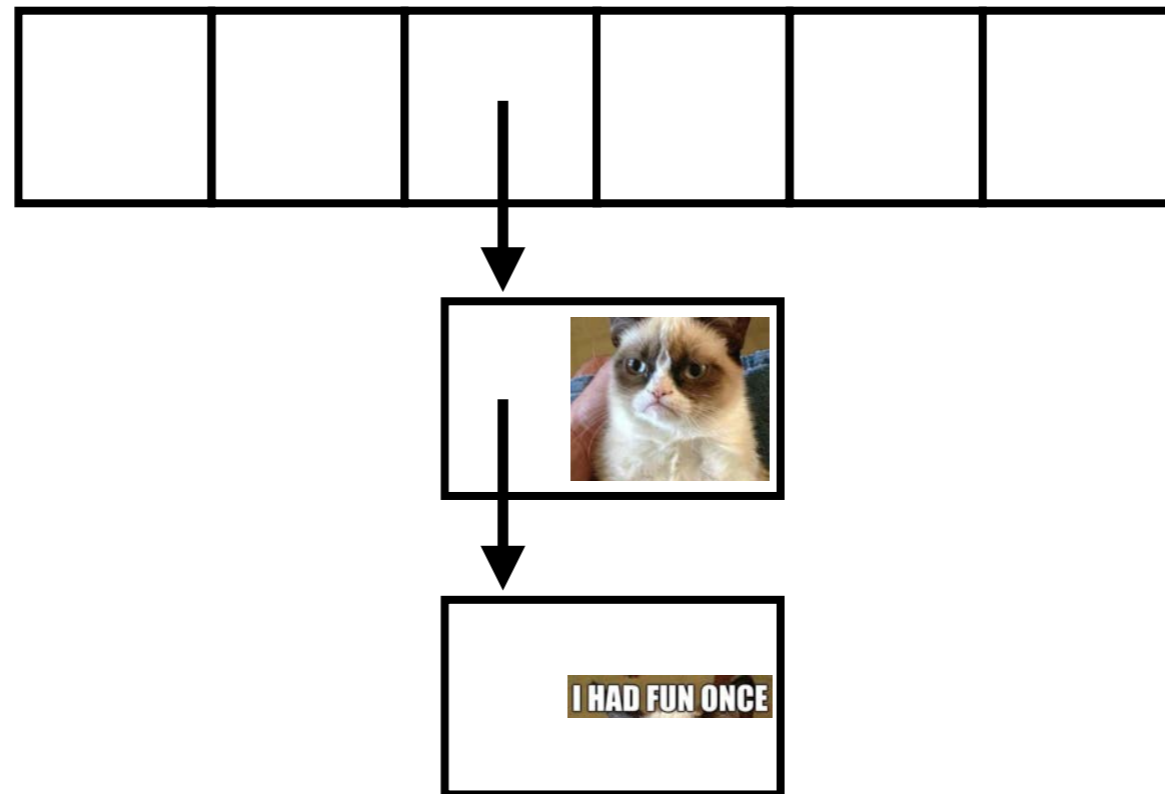
`i = hash_function()`



Dedup

Compression via deduplication


hash_function()



Dedup

Compression via deduplication

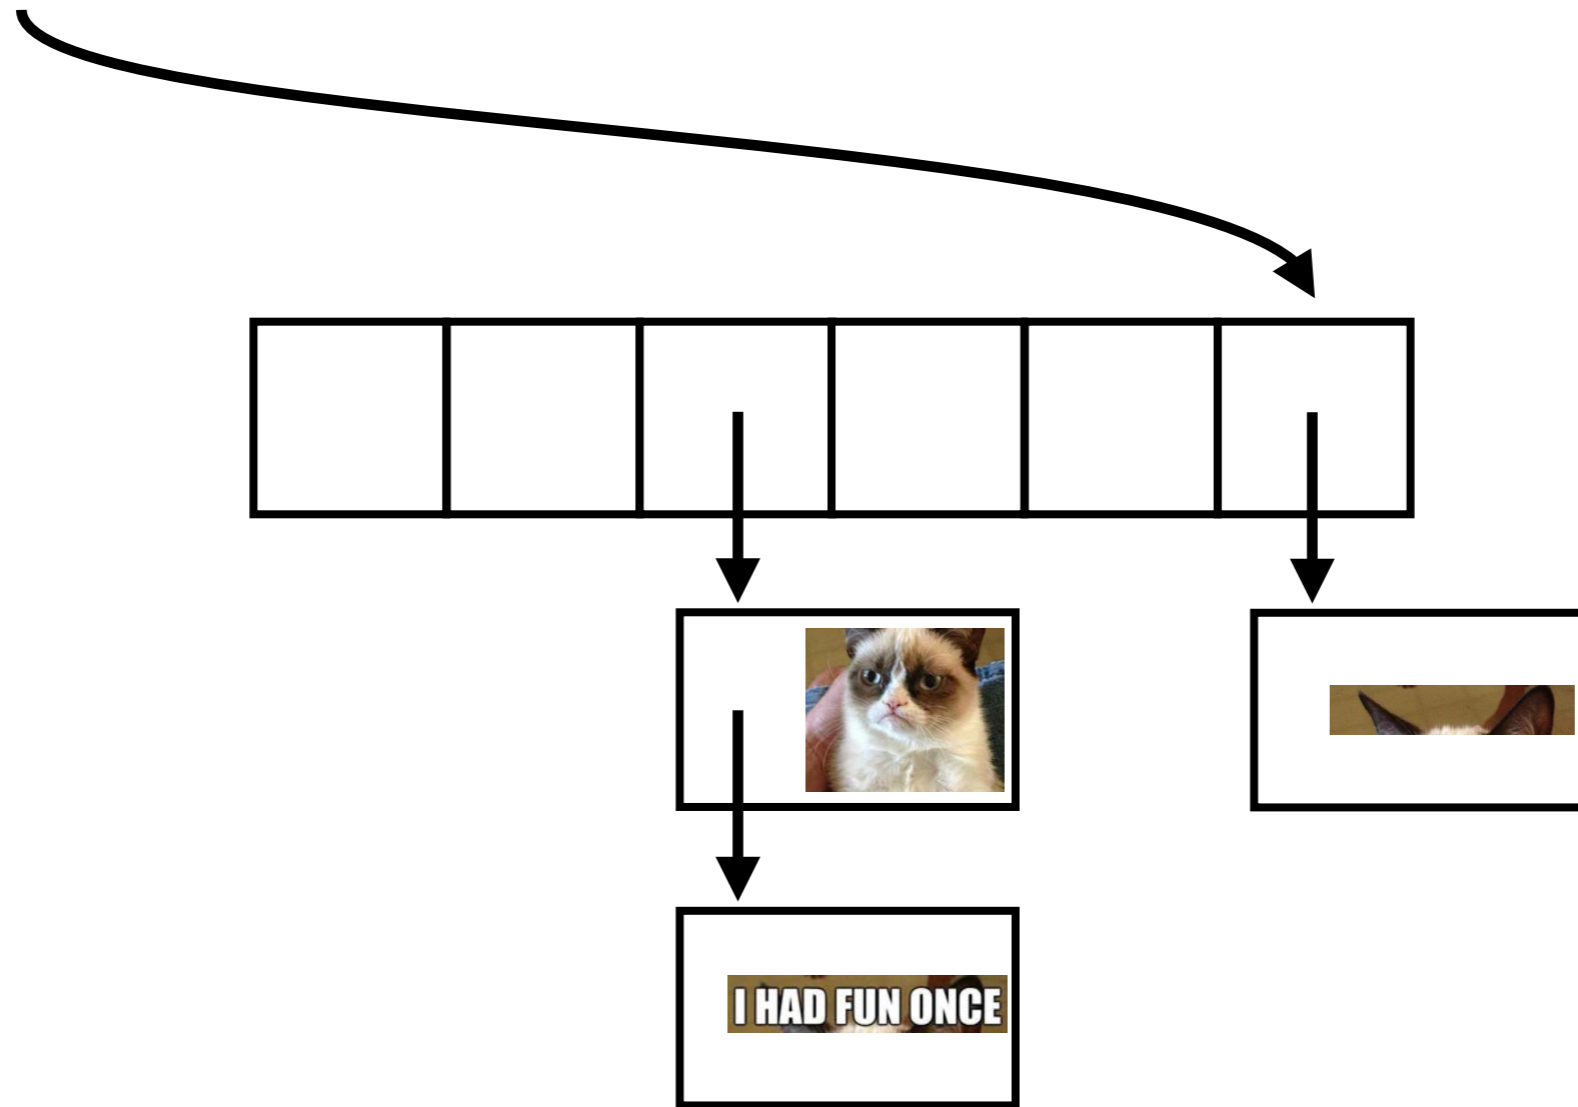
$i = \text{hash_function}(\text{img})$



Dedup

Compression via deduplication

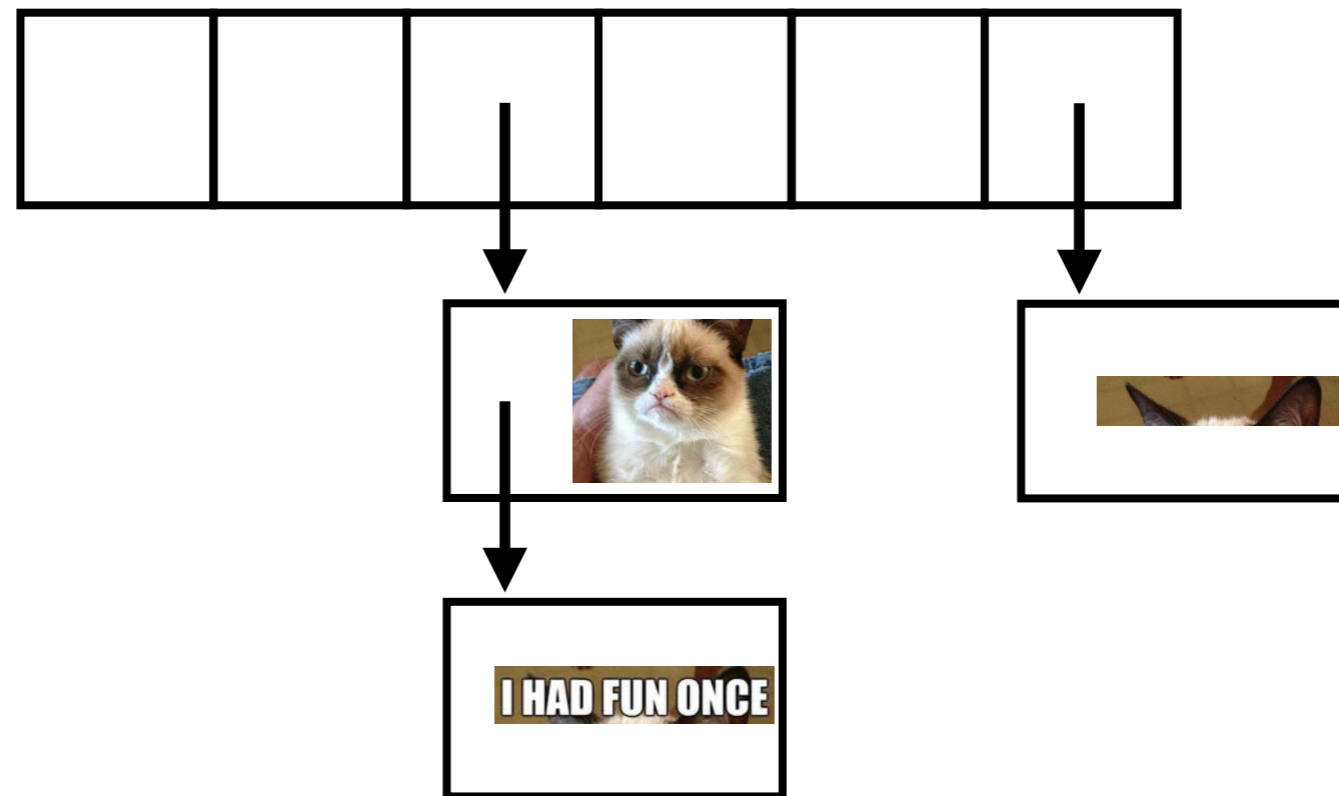
`i = hash_function()`



Dedup

Compression via deduplication


Hash table is accessed concurrently by many threads



Dedup

Compression via deduplication

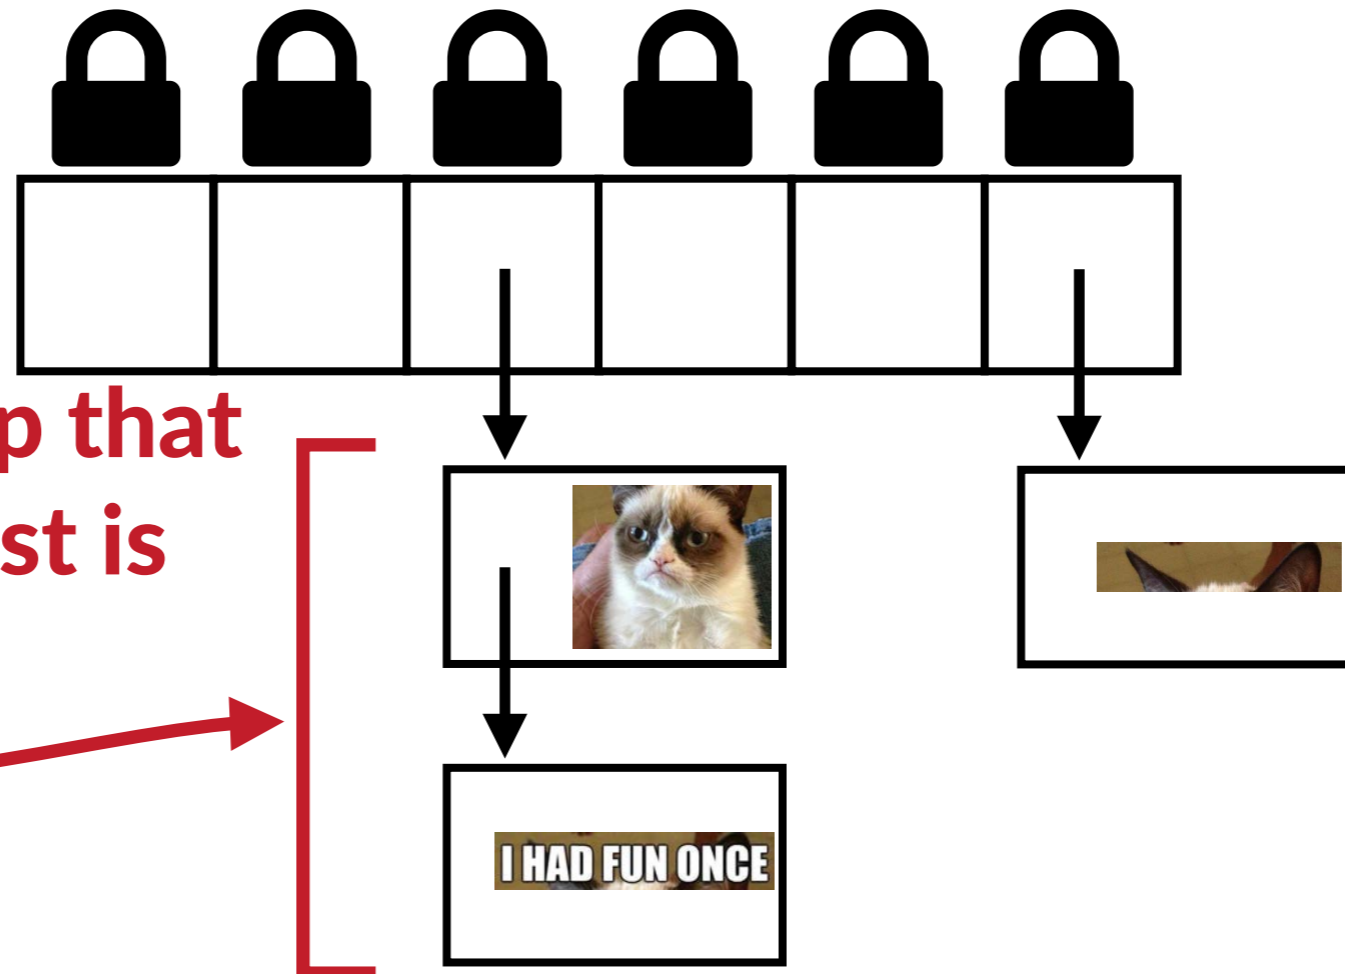
Hash table is accessed concurrently by many threads



Dedup

Compression via deduplication

Hash table is accessed concurrently by many threads




Coz says the loop that accesses this list is important



Dedup


Compression via deduplication



Dedup

Compression via deduplication

More hash buckets should lead to fewer collisions




Dedup

Compression via deduplication

More hash buckets should lead to fewer collisions

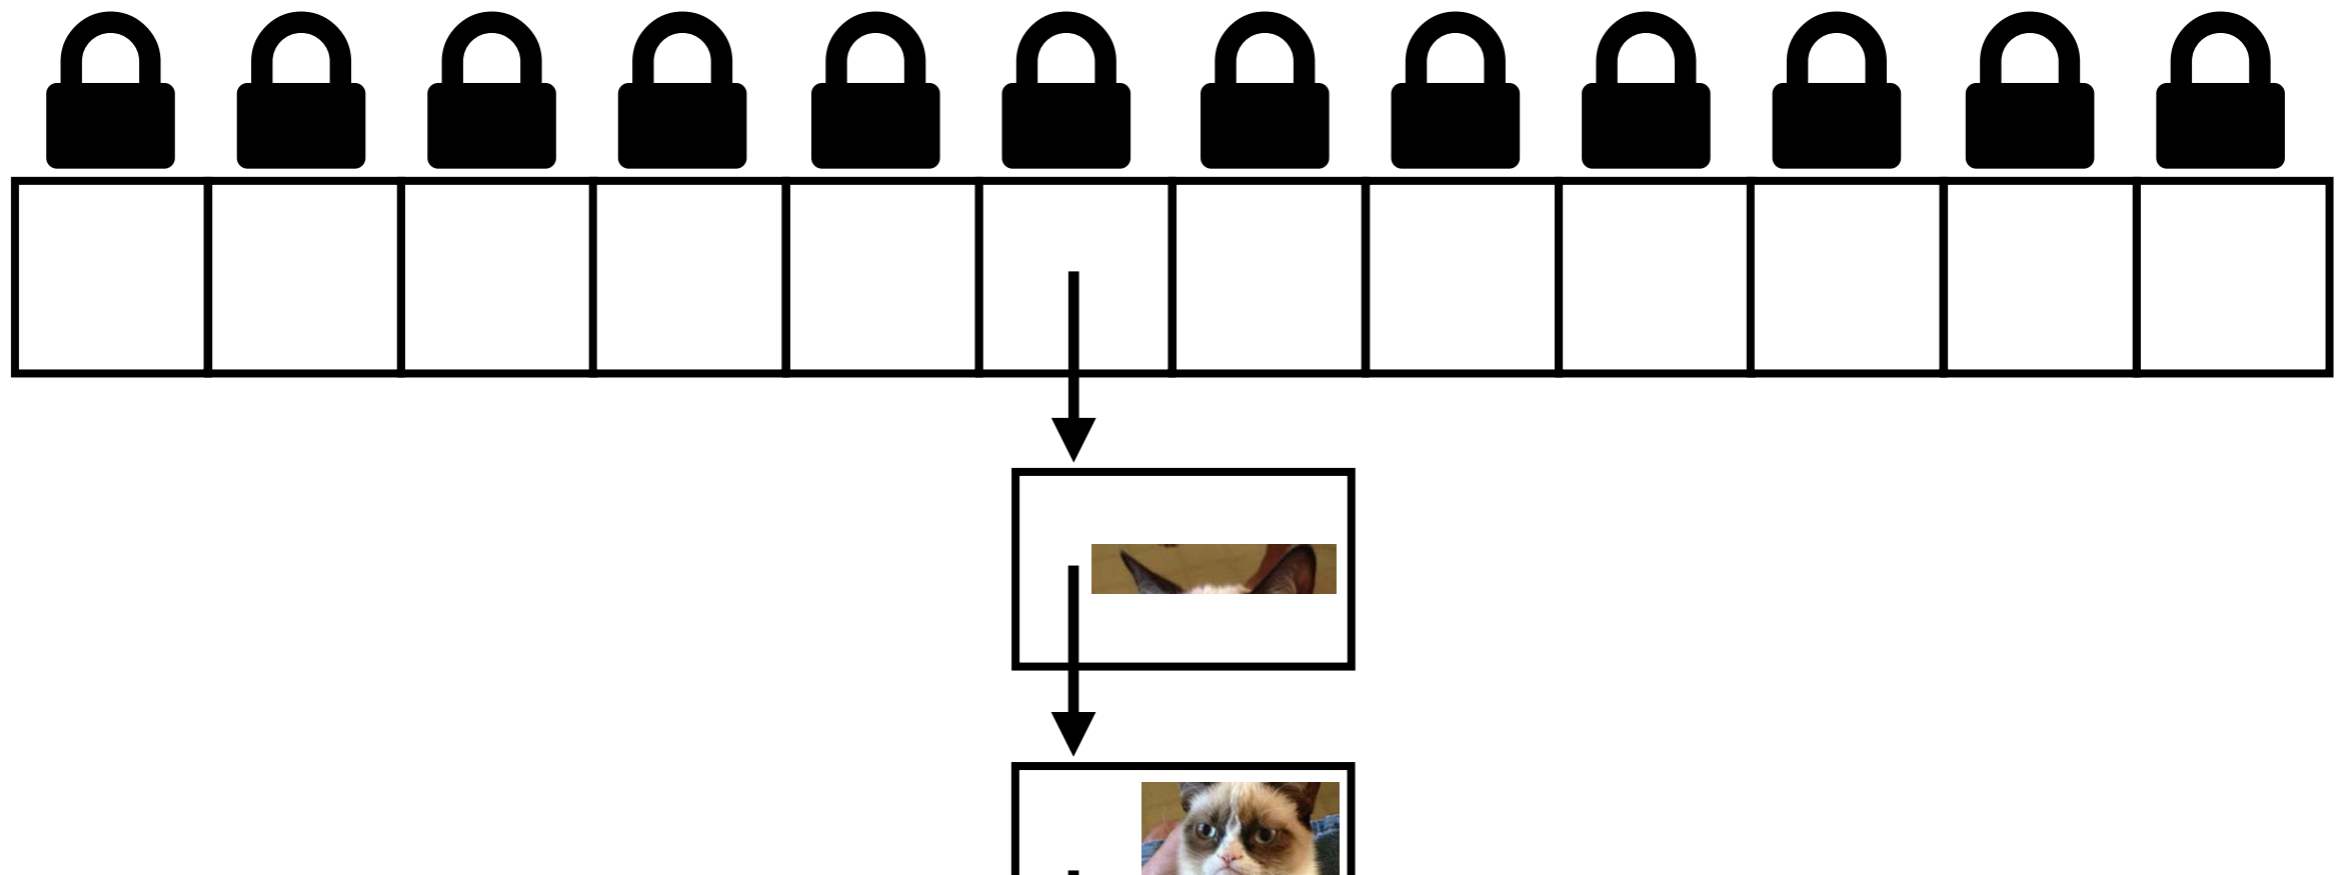
No performance improvement



Dedup

Compression via deduplication

What else could be causing collisions?

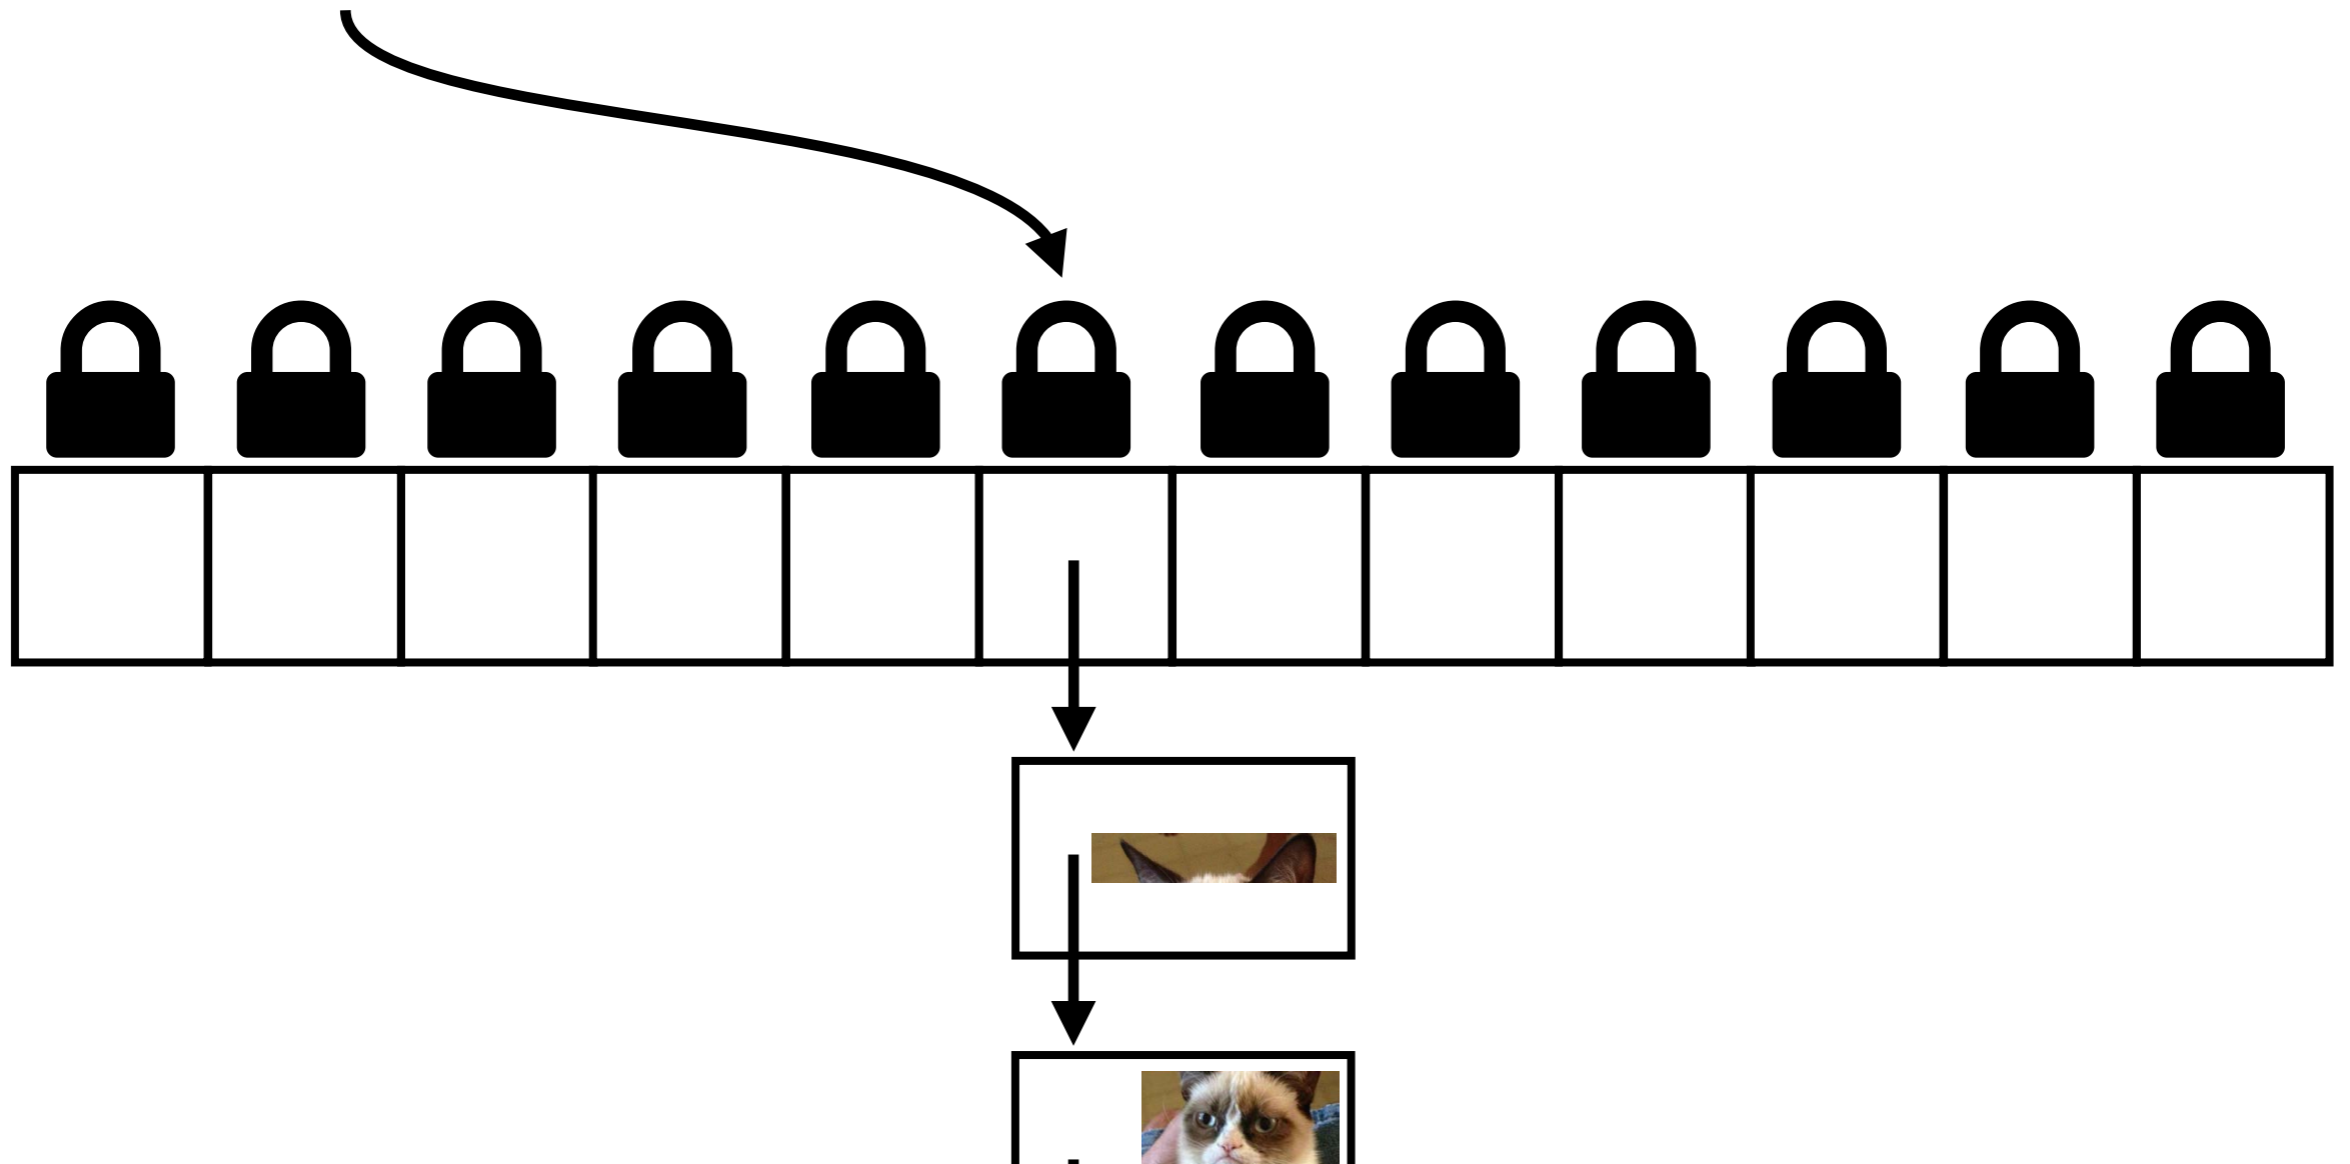


Dedup

Compression via deduplication

What else could be causing collisions?


$i = \text{hash_function}(\text{img})$



Dedup

Compression via deduplication


Horrible hash function!



Dedup

Compression via deduplication

Horrible hash function!




Bin Utilization

2.3%

Dedup

Compression via deduplication




Bin Utilization

2.3%

Dedup

Compression via deduplication



Bin Utilization

2.3%

82%

Dedup

Compression via deduplication



Bin Utilization

2.3%


9% Speedup

82%

Dedup

Compression via deduplication

What did Coz predict?




Dedup

Compression via deduplication

What did Coz predict?

Blocks per-bucket




Dedup

Compression via deduplication

What did Coz predict?

Blocks per-bucket

Before: 76.7



Dedup


Compression via deduplication

What did Coz predict?

Blocks per-bucket

Before: 76.7

After: 2.09



Dedup

Compression via deduplication


What did Coz predict?

Blocks per-bucket

Before: 76.7

After: 2.09

96% traversal speedup



Dedup

Compression via deduplication

What did Coz predict?


Blocks per-bucket

Before: 76.7

After: 2.09

96% traversal speedup

**9% predicted speedup,
exactly what we observed**

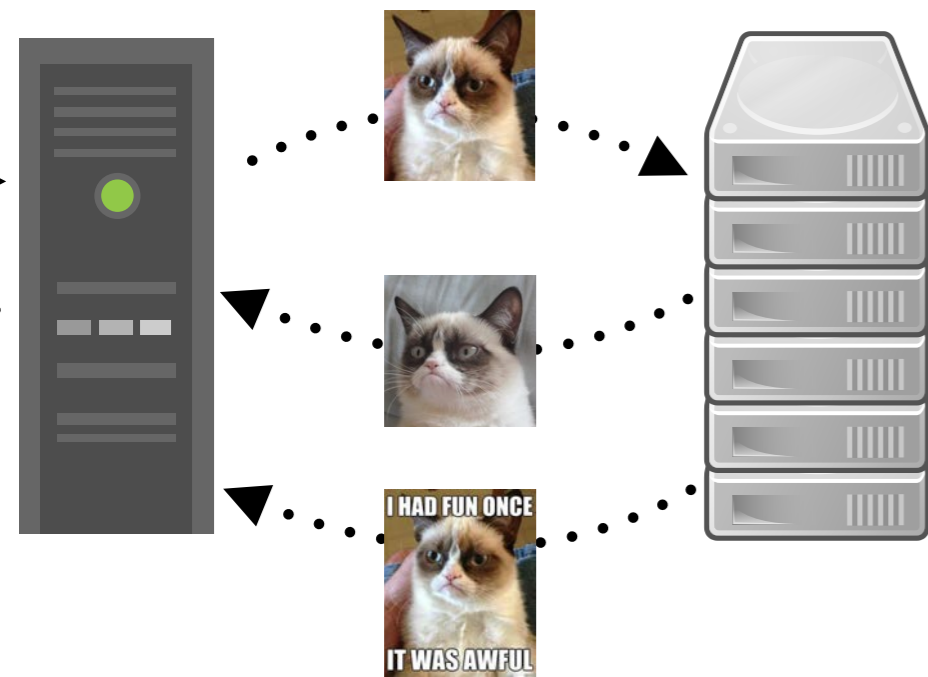


Using Causal Profiling on Ogle



dedup
compression

ferret
image comparison




Using Causal Profiling on Ogle

 SQLite 25%

dedup
compression

ferret
image comparison

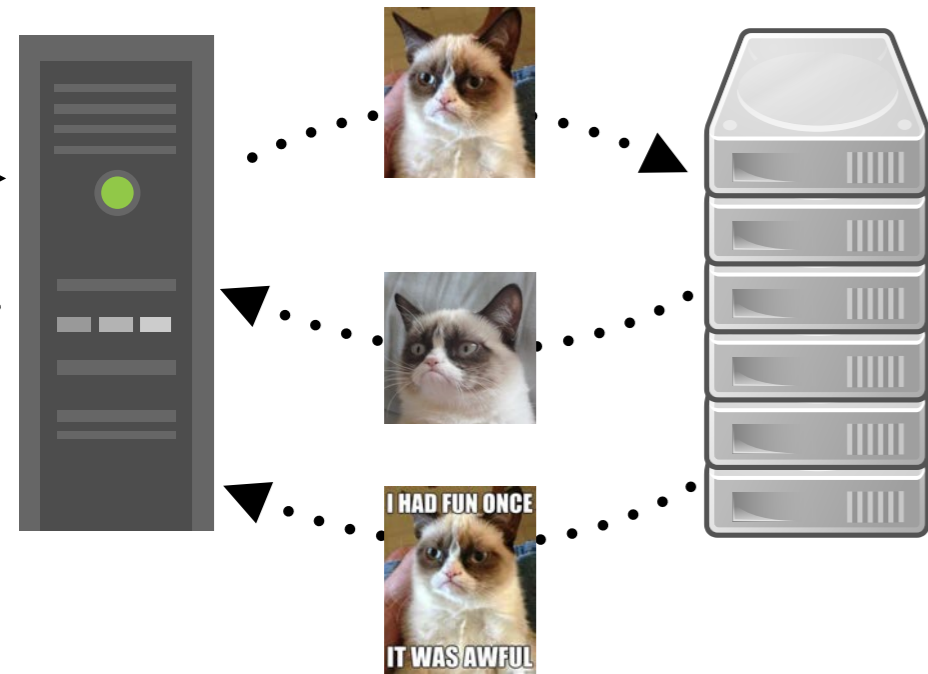


Using Causal Profiling on Ogle

 SQLite 25%

dedup
compression 9%

ferret
image comparison




Using Causal Profiling on Ogle

 SQLite 25%

dedup
compression 9%

ferret
image comparison 21%



Summary of Optimizations

Benchmark	Speedup	Diff Size	Change Summary
memcached	9.39%	-6, +2	removed unnecessary locks
sqlite	25.60%	-3, +3	removed DIY vtable implementation
blackscholes	2.56%	-61, +4	manual common subexpression elimination
dedup	8.95%	-3, +3	fixed degenerate hash function
ferret	21.27%	-4, +4	rebalanced pipeline thread allocation
fluidanimate	37.50%	-1, +0	removed custom barrier with high contention
streamcluster	68.40%	-1, +0	removed custom barrier with high contention
swaptions	15.80%	-10, +16	reordered loop nests







Effective Performance Profiling





COZ






```
% sudo apt install coz-profiler
```

































python



python




python

the "C"




python



the "C"

SCALENE


SCRIPTING-LANGUAGE
AWARE
PROFILING




SCALENE

github.com/emeryberger/scalene


```
% scalene yourprogram.py
```




```
% scalene --cpu-only yourprogram.py
```



```
% scalene --cpu-only --html yourprogram.py
```




```
% scalene --help
```






```
#!/usr/bin/env python3
import numpy as np

def main():
    x = np.array(range(10**7))
    y = np.array(np.random.uniform(0, 100, size=(10**8)))


main()
```






Memory usage:  (max: 1618.94MB)
 test2-2.py: % of time = 100.00% out of 6.35s.

Line	CPU % Python	CPU % native	Sys %	Mem % Python	Net (MB)	Memory usage over time / %	Copy (MB/s)	test2-2.py
1								<code>#!/usr/bin/env python3</code>
2	4%	1%	26%	97%	9		3	<code>import numpy as np</code>
3								
4								<code>def main():</code>
5		60%		93%	84			<code> x = np.array(range(10**7))</code>
6	1%	34%			763	 87%	120	<code> y = np.array(np.random.uniform(0, 100, size=(10**8)))</code>
7								
8								<code>main()</code>

generated by the [scalene](#) profiler





Memory usage:  (max: 1618.94MB)
 test2-2.py: % of time = 100.00% out of 6.35s.



Line	CPU % Python	CPU % native	Sys %	Mem % Python	Net (MB)	Memory usage over time / %	Copy (MB/s)	test2-2.py
1								<code>#!/usr/bin/env python3</code>
2	4%	1%	26%	97%	9		3	<code>import numpy as np</code>
3								
4								<code>def main():</code>
5		60%		93%	84			<code> x = np.array(range(10**7))</code>
6	1%	34%			763	 87%	120	<code> y = np.array(np.random.uniform(0, 100, size=(10**8)))</code>
7								
8								<code>main()</code>

generated by the [scalene](#) profiler

**CPU
PYTHON
vs. NATIVE
+ SYS%**




Memory usage:  (max: 1618.94MB)
 test2-2.py: % of time = 100.00% out of 6.35s.

Line	CPU % Python	CPU % native	Sys %	Mem % Python	Net (MB)	Memory usage over time / %	Copy (MB/s)	test2-2.py
1								<code>#!/usr/bin/env python3</code>
2	4%	1%	26%	97%	9		3	<code>import numpy as np</code>
3								
4								<code>def main():</code>
5		60%		93%	84			<code> x = np.array(range(10**7))</code>
6	1%	34%			763	 87%	120	<code> y = np.array(np.random.uniform(0, 100, size=(10**8)))</code>
7								
8								<code>main()</code>


generated by the [scalene](#) profiler

CPU
PYTHON
 vs. NATIVE
 + SYS%

MEMORY
PYTHON
 vs. NATIVE



MEMORY USAGE OVER TIME

Memory usage:  (max: 1618.94MB)
 test2-2.py: % of time = 100.00% out of 6.35s.

Line	CPU % Python	CPU % native	Sys %	Mem % Python	Net (MB)	Memory usage over time / %	Copy (MB/s)	test2-2.py
1								<code>#!/usr/bin/env python3</code>
2	4%	1%	26%	97%	9		3	<code>import numpy as np</code>
3								
4								<code>def main():</code>
5		60%		93%	84			<code> x = np.array(range(10**7))</code>
6	1%	34%			763	87%	120	<code> y = np.array(np.random.uniform(0, 100, size=(10**8)))</code>
7								
8								<code>main()</code>


generated by the [scalene](#) profiler

**CPU
PYTHON
vs. NATIVE
+ SYS%**

**MEMORY
PYTHON
vs. NATIVE**

**MEMORY
USAGE
OVER TIME**

**% OF MEM
ALLOCATED**



MEMORY USAGE OVER TIME

Memory usage: (max: 1618.94MB)
 test2-2.py: % of time = 100.00% out of 6.35s.

Line	CPU % Python	CPU % native	Sys %	Mem % Python	Net (MB)	Memory usage over time / %	Copy (MB/s)	test2-2.py
1								<code>#!/usr/bin/env python3</code>
2	4%	1%	26%	97%	9		3	<code>import numpy as np</code>
3								
4								<code>def main():</code>
5		60%		93%	84			<code> x = np.array(range(10**7))</code>
6	1%	34%			763	87%	120	<code> y = np.array(np.random.uniform(0, 100, size=(10**8)))</code>
7								
8								<code>main()</code>

generated by the [scalene](#) profiler


**CPU
PYTHON
vs. NATIVE
+ SYS%**

**MEMORY
PYTHON
vs. NATIVE**



**MEMORY
USAGE
OVER
TIME

% OF MEM
ALLOCATED**

**COPY
VOLUME
(MB/s)**





Memory usage:  (max: 1618.94MB)
 test2-2.py: % of time = 100.00% out of 6.35s.

Line	CPU % Python	CPU % native	Sys %	Mem % Python	Net (MB)	Memory usage over time / %	Copy (MB/s)	test2-2.py
1								<code>#!/usr/bin/env python3</code>
2	4%	1%	26%	97%	9		3	<code>import numpy as np</code>
3								
4								<code>def main():</code>
5		60%		93%	84			<code> x = np.array(range(10**7))</code>
6	1%	34%			763	 87%	120	<code> y = np.array(np.random.uniform(0, 100, size=(10**8)))</code>
7								
8								<code>main()</code>

generated by the [scalene](#) profiler



Memory usage:  (max: 1618.94MB)
test2-2.py: % of time = 100.00% out of 6.35s.

Line	CPU % Python	CPU % native	Sys %	Mem % Python	Net (MB)	Memory usage over time / %	Copy (MB/s)	test2-2.py
1								<code>#!/usr/bin/env python3</code>
2	4%	1%	26%	97%	9		3	<code>import numpy as np</code>
3								
4								<code>def main():</code>
5		60%		93%	84			<code> x = np.array(range(10**7))</code>
6	1%	34%			763	 87%	120	<code> y = np.array(np.random.uniform(0, 100, size=(10**8)))</code>
7								
8								<code>main()</code>

generated by the [scalene](#) profiler

34% of runtime in native code

Memory usage:  (max: 1618.94MB)
test2-2.py: % of time = 100.00% out of 6.35s.



Line	CPU % Python	CPU % native	Sys %	Mem % Python	Net (MB)	Memory usage over time / %	Copy (MB/s)	test2-2.py
1								<code>#!/usr/bin/env python3</code>
2	4%	1%	26%	97%	9		3	<code>import numpy as np</code>
3								
4								<code>def main():</code>
5		60%		93%	84			<code> x = np.array(range(10**7))</code>
6	1%	34%			763	 87%	120	<code> y = np.array(np.random.uniform(0, 100, size=(10**8)))</code>
7								
8								<code>main()</code>

generated by the [scalene](#) profiler

34% of runtime in native code

763MB allocated in native code

Memory usage:  (max: 1618.94MB)
test2-2.py: % of time = 100.00% out of 6.35s.

Line	CPU % Python	CPU % native	Sys %	Mem % Python	Net (MB)	Memory usage over time / %	Copy (MB/s)	test2-2.py
1								<code>#!/usr/bin/env python3</code>
2	4%	1%	26%	97%	9		3	<code>import numpy as np</code>
3								
4								<code>def main():</code>
5		60%		93%	84			<code> x = np.array(range(10**7))</code>
6	1%	34%			763	 87%	120	<code> y = np.array(np.random.uniform(0, 100, size=(10**8)))</code>
7								
8								<code>main()</code>



generated by the [scalene](#) profiler

34% of runtime in native code

763MB allocated in native code

87% of memory activity

Memory usage:  (max: 1618.94MB)
test2-2.py: % of time = 100.00% out of 6.35s.

Line	CPU % Python	CPU % native	Sys %	Mem % Python	Net (MB)	Memory usage over time / %	Copy (MB/s)	test2-2.py
1								<code>#!/usr/bin/env python3</code>
2	4%	1%	26%	97%	9		3	<code>import numpy as np</code>
3								
4								<code>def main():</code>
5		60%		93%	84			<code> x = np.array(range(10**7))</code>
6	1%	34%			763	 87%	120	<code> y = np.array(np.random.uniform(0, 100, size=(10**8)))</code>
7								
8								<code>main()</code>



generated by the [scalene](#) profiler

34% of runtime in native code

763MB allocated in native code

87% of memory activity "sawtooth" pattern

Memory usage:  (max: 1618.94MB)
test2-2.py: % of time = 100.00% out of 6.35s.

Line	CPU % Python	CPU % native	Sys %	Mem % Python	Net (MB)	Memory usage over time / %	Copy (MB/s)	test2-2.py
1								<code>#!/usr/bin/env python3</code>
2	4%	1%	26%	97%	9		3	<code>import numpy as np</code>
3								
4								<code>def main():</code>
5		60%		93%	84			<code> x = np.array(range(10**7))</code>
6	1%	34%			763	 87%	120	<code> y = np.array(np.random.uniform(0, 100, size=(10**8)))</code>
7								
8								<code>main()</code>

generated by the [scalene](#) profiler



34% of runtime in native code

763MB allocated in native code

87% of memory activity "sawtooth" pattern

120MB/s copying

Memory usage:  (max: 1618.94MB)
test2-2.py: % of time = 100.00% out of 6.35s.

Line	CPU % Python	CPU % native	Sys %	Mem % Python	Net (MB)	Memory usage over time / %	Copy (MB/s)	test2-2.py
1								<code>#!/usr/bin/env python3</code>
2	4%	1%	26%	97%	9		3	<code>import numpy as np</code>
3								
4								
5		60%		93%	84			<code>def main():</code>
6	1%	34%			763	 87%	120	<code> x = np.array(range(10**7))</code>
7								<code> y = np.array(np.random.uniform(0, 100, size=(10**8)))</code>
8								<code>main()</code>

generated by the [scalene](#) profiler

34% of runtime in native code



763MB allocated in native code

87% of memory activity "sawtooth" pattern

120MB/s copying

converts to numpy array

Memory usage:  (max: 1618.94MB)
test2-2.py: % of time = 100.00% out of 6.35s.

Line	CPU % Python	CPU % native	Sys %	Mem % Python	Net (MB)	Memory usage over time / %	Copy (MB/s)	test2-2.py
1								<code>#!/usr/bin/env python3</code>
2	4%	1%	26%	97%	9		3	<code>import numpy as np</code>
3								
4								<code>def main():</code>
5		60%		93%	84			<code> x = np.array(range(10**7))</code>
6	1%	34%			763	 87%	120	<code> y = np.array(np.random.uniform(0, 100, size=(10**8)))</code>
7								
8								<code>main()</code>

generated by the [scalene](#) profiler


34% of runtime in native code



763MB allocated in native code

87% of memory activity "sawtooth" pattern


120MB/s copying



converts to numpy array
already a numpy array!

Memory usage:  (max: 1618.94MB)
test2-2.py: % of time = 100.00% out of 6.35s.


Line	CPU % Python	CPU % native	Sys %	Mem % Python	Net (MB)	Memory usage over time / %	Copy (MB/s)	test2-2.py
1								<code>#!/usr/bin/env python3</code>
2	4%	1%	26%	97%	9		3	<code>import numpy as np</code>
3								
4								<code>def main():</code>
5		60%		93%	84			<code> x = np.array(range(10**7))</code>
6	1%	34%			763	 87%	120	<code> y = np.array(np.random.uniform(0, 100, size=(10**8)))</code>
7								
8								<code>main()</code>


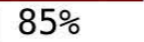
generated by the [scalene](#) profiler

Memory usage:  (max: 1618.94MB)
 test2-2.py: % of time = 100.00% out of 6.35s.


Line	CPU % Python	CPU % native	Sys %	Mem % Python	Net (MB)	Memory usage over time / %	Copy (MB/s)	test2-2.py
1								<code>#!/usr/bin/env python3</code>
2	4%	1%	26%	97%	9		3	<code>import numpy as np</code>
3								
4								<code>def main():</code>
5		60%		93%	84			<code> x = np.array(range(10**7))</code>
6	1%	34%			763	 87%	120	<code> y = np.array(np.random.uniform(0, 100, size=(10**8)))</code>
7								
8								<code>main()</code>



generated by the [scalene](#) profiler

Memory usage:  (max: 865.97MB)
 test2-3.py: % of time = 100.00% out of 5.40s.


Line	CPU % Python	CPU % native	Sys %	Mem % Python	Net (MB)	Memory usage over time / %	Copy (MB/s)	test2-3.py
1								<code>#!/usr/bin/env python3</code>
2	6%	1%	34%	93%	21		3	<code>import numpy as np</code>
3								
4								<code>#@profile</code>
5								<code>def main():</code>
6		67%		94%	82			<code> x = np.array(range(10**7))</code>
7		25%	5%		763	 85%		<code> y = np.random.uniform(0, 100, size=(10**8))</code>
8								
9								<code>main()</code>


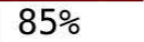
generated by the [scalene](#) profiler

Memory usage:  (max: 1618.94MB)
 test2-2.py: % of time = 100.00% out of 6.35s.


Line	CPU % Python	CPU % native	Sys %	Mem % Python	Net (MB)	Memory usage over time / %	Copy (MB/s)	test2-2.py
1								<code>#!/usr/bin/env python3</code>
2	4%	1%	26%	97%	9		3	<code>import numpy as np</code>
3								
4								<code>def main():</code>
5		60%		93%	84			<code> x = np.array(range(10**7))</code>
6	1%	34%			763	 87%	120	<code> y = np.array(np.random.uniform(0, 100, size=(10**8)))</code>
7								
8								<code>main()</code>



generated by the [scalene](#) profiler

Memory usage:  (max: 865.97MB)
 test2-3.py: % of time = 100.00% out of 5.40s.


Line	CPU % Python	CPU % native	Sys %	Mem % Python	Net (MB)	Memory usage over time / %	Copy (MB/s)	test2-3.py
1								<code>#!/usr/bin/env python3</code>
2	6%	1%	34%	93%	21		3	<code>import numpy as np</code>
3								
4								<code>#@profile</code>
5								<code>def main():</code>
6		67%		94%	82			<code> x = np.array(range(10**7))</code>
7		25%	5%		763	 85%		<code> y = np.random.uniform(0, 100, size=(10**8))</code>
8								
9								<code>main()</code>


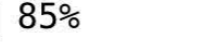
generated by the [scalene](#) profiler

Memory usage:  (max: 1618.94MB)
 test2-2.py: % of time = 100.00% out of 6.35s.

Line	CPU % Python	CPU % native	Sys %	Mem % Python	Net (MB)	Memory usage over time / %	Copy (MB/s)	test2-2.py
1								<code>#!/usr/bin/env python3</code>
2	4%	1%	26%	97%	9		3	<code>import numpy as np</code>
3								
4								<code>def main():</code>
5		60%		93%	84			<code>x = np.array(range(10**7))</code>
6	1%	34%			763	 87%	120	<code>y = np.array(np.random.uniform(0, 100, size=(10**8)))</code>
7								
8								<code>main()</code>



generated by the [scalene](#) profiler

Memory usage:  (max: 865.97MB)
 test2-3.py: % of time = 100.00% out of 5.40s.


Line	CPU % Python	CPU % native	Sys %	Mem % Python	Net (MB)	Memory usage over time / %	Copy (MB/s)	test2-3.py
1								<code>#!/usr/bin/env python3</code>
2	6%	1%	34%	93%	21		3	<code>import numpy as np</code>
3								
4								<code>#@profile</code>
5								<code>def main():</code>
6		67%		94%	82			<code>x = np.array(range(10**7))</code>
7		25%	5%		763	 85%		<code>y = np.random.uniform(0, 100, size=(10**8))</code>
8								
9								<code>main()</code>



generated by the [scalene](#) profiler

Memory usage:  (max: 1618.94MB)
 test2-2.py: % of time = 100.00% out of 6.35s.

Line	CPU % Python	CPU % native	Sys %	Mem % Python	Net (MB)	Memory usage over time / %	Copy (MB/s)	test2-2.py
1								<code>#!/usr/bin/env python3</code>
2	4%	1%	26%	97%	9		3	<code>import numpy as np</code>
3								
4								<code>def main():</code>
5		60%		93%	84			<code>x = np.array(range(10**7))</code>
6	1%	34%			763	 87%	120	<code>y = np.array(np.random.uniform(0, 100, size=(10**8)))</code>
7								
8								<code>main()</code>



generated by the [scalene](#) profiler

Memory usage:  (max: 865.97MB)
 test2-3.py: % of time = 100.00% out of 5.40s.

Line	CPU % Python	CPU % native	Sys %	Mem % Python	Net (MB)	Memory usage over time / %	Copy (MB/s)	test2-3.py
1								<code>#!/usr/bin/env python3</code>
2	6%	1%	34%	93%	21		3	<code>import numpy as np</code>
3								
4								<code>#@profile</code>
5								<code>def main():</code>
6		67%		94%	82			<code>x = np.array(range(10**7))</code>
7		25%	5%		763	 85%		<code>y = np.random.uniform(0, 100, size=(10**8))</code>
8								
9								<code>main()</code>



generated by the [scalene](#) profiler

Memory usage:  (max: 1618.94MB)
 test2-2.py: % of time = 100.00% out of 6.35s.


Line	CPU % Python	CPU % native	Sys %	Mem % Python	Net (MB)	Memory usage over time / %	Copy (MB/s)	test2-2.py
1								<code>#!/usr/bin/env python3</code>
2	4%	1%	26%	97%	9		3	<code>import numpy as np</code>
3								
4								<code>def main():</code>
5		60%		93%	84			<code>x = np.array(range(10**7))</code>
6	1%	34%			763	 87%	120	<code>y = np.array(np.random.uniform(0, 100, size=(10**8)))</code>
7								
8								<code>main()</code>



generated by the [scalene](#) profiler

Memory usage:  (max: 865.97MB)
 test2-3.py: % of time = 100.00% out of 5.46s.


Line	CPU % Python	CPU % native	Sys %	Mem % Python	Net (MB)	Memory usage over time / %	Copy (MB/s)	test2-3.py
1								<code>#!/usr/bin/env python3</code>
2	6%	1%	34%	93%	21		3	<code>import numpy as np</code>
3								
4								<code>#@profile</code>
5								<code>def main():</code>
6		67%		94%	82			<code>x = np.array(range(10**7))</code>
7		25%	5%		763	 85%		<code>y = np.random.uniform(0, 100, size=(10**8))</code>
8								
9								<code>main()</code>



generated by the [scalene](#) profiler

Memory usage:  (max: 160.94MB)
 test2-2.py: % of time = 100.00% out of 6.35s.

Line	CPU % Python	CPU % native	Sys %	Mem % Python	Net (MB)	Memory usage over time / %	Copy (MB/s)	test2-2.py
1								<code>#!/usr/bin/env python3</code>
2	4%	1%	26%	97%	9		3	<code>import numpy as np</code>
3								
4								<code>def main():</code>
5		60%		93%	84			<code>x = np.array(range(10**7))</code>
6	1%	34%			763	 87%	120	<code>y = np.array(np.random.uniform(0, 100, size=(10**8)))</code>
7								
8								<code>main()</code>

generated by the [scalene](#) profiler


Memory usage:  (max: 865.97MB)
 test2-3.py: % of time = 100.00% out of 5.40s.

Line	CPU % Python	CPU % native	Sys %	Mem % Python	Net (MB)	Memory usage over time / %	Copy (MB/s)	test2-3.py
1								<code>#!/usr/bin/env python3</code>
2	6%	1%	34%	93%	21		3	<code>import numpy as np</code>
3								
4								<code>#@profile</code>
5								<code>def main():</code>
6		67%		94%	82			<code>x = np.array(range(10**7))</code>
7		25%	5%		763	 85%		<code>y = np.random.uniform(0, 100, size=(10**8))</code>
8								
9								<code>main()</code>

generated by the [scalene](#) profiler

SCALENE

SCRIPTING-LANGUAGE
AWARE
PROFILING




precise

<u>CPU</u> PYTHON vs. NATIVE + SYS%	<u>MEMORY</u> PYTHON vs. NATIVE	<u>MEMORY</u> USAGE OVER TIME % OF MEM ALLOCATED	<u>COPY</u> VOLUME (MB/s)
--	---------------------------------------	---	---------------------------------

SCALENE

SCRIPTING-LANGUAGE
AWARE
PROFILING





SCALENE

precise


fast

<u>CPU</u> PYTHON vs. NATIVE + SYS%	<u>MEMORY</u> PYTHON vs. NATIVE	<u>MEMORY</u> USAGE OVER TIME % OF MEM ALLOCATED	<u>COPY</u> VOLUME (MB/s)
--	---------------------------------------	---	---------------------------------


Normalized profiler execution time (mdp)




Normalized profiler execution time (mdp)




memory_profiler: 152x
(12m38s vs. 5s)




Normalized profiler execution time (mdp)



memory_profiler: 152x
(12m38s vs. 5s)



DEFERRED SIGNAL DELIVERY




python™






DEFERRED SIGNAL DELIVERY



python™




python™







DEFERRED SIGNAL DELIVERY



python™



python™




python™







DEFERRED SIGNAL DELIVERY




python™



python™



python™




python™







DEFERRED SIGNAL DELIVERY




python™




python™



python™



python™




python™







DEFERRED SIGNAL DELIVERY




python™




python™



python™



python™




python™







DEFERRED SIGNAL DELIVERY




python™




python™



python™



python™




python™







DEFERRED SIGNAL DELIVERY




python™




python™




python™



python™



python™




python™







DEFERRED SIGNAL DELIVERY




python™




python™




python™



python™



python™



python™




INFERRING EXECUTION TIME




(wall clock)

INFERRING EXECUTION TIME



(wall clock)

INFERRING EXECUTION TIME



(wall clock)

```
python_time += interval  
c_time += delay - interval
```

DEFERRED SIGNAL DELIVERY



delay



(wall clock)

```
python_time += interval  
c_time += delay - interval
```



(virtual time)

```
sys_time = 1 - elapsed_virtual /  
elapsed_wallclock
```

DEFERRED SIGNAL DELIVERY



delay



(wall clock)

```
python_time += interval
c_time += delay - interval
```

takesometime.py: % of time = 100.00% out of 1.03s.

Line	CPU % Python	CPU % native	Sys %	takesometime.py
1				<code>import time</code>
2				<code>import taketime</code>
3				
4				<code>c_its_per_sec = 250000</code>
5				<code>python_its_per_sec = 1080000</code>
6				
7				<code>import timeit</code>
8				
9				<code>def py_taketime(n):</code>
10				<code> d = 1.01</code>
11				<code> for i in range(n):</code>
12	7%	1%		<code> for j in range(10):</code>
13	18%	3%		<code> d = d * d + d * d</code>
14				<code> return n</code>
15				
16				<code>c_secs = .7</code>
17				<code>python_secs = .3</code>
18				
19				<code>start = time.process_time()</code>
20	1%	70%		<code>p = taketime.run(int(c_secs * c_its_per_sec))</code>
21				<code>q = py_taketime(int(python_secs * python_its_per_sec))</code>
22				<code>end = time.process_time()</code>
23				<code>print(end-start)</code>

DEFERRED SIGNAL DELIVERY



delay



(wall clock)

```
python_time += interval
c_time += delay - interval
```

takesometime.py: % of time = 100.00% out of 1.03s.

≈30%

≈70%

Line	CPU % Python	CPU % native	Sys %	takesometime.py
1				<code>import time</code>
2				<code>import taketime</code>
3				
4				<code>c_its_per_sec = 250000</code>
5				<code>python_its_per_sec = 1080000</code>
6				
7				<code>import timeit</code>
8				
9				<code>def py_taketime(n):</code>
10				<code> d = 1.01</code>
11				<code> for i in range(n):</code>
12	7%	1%		<code> for j in range(10):</code>
13	18%	3%		<code> d = d * d + d * d</code>
14				<code> return n</code>
15				
16				<code>c_secs = .7</code>
17				<code>python_secs = .3</code>
18				
19				<code>start = time.process_time()</code>
20	1%	70%		<code>p = taketime.run(int(c_secs * c_its_per_sec))</code>
21				<code>q = py_taketime(int(python_secs * python_its_per_sec))</code>
22				<code>end = time.process_time()</code>
23				<code>print(end-start)</code>

DEFERRED SIGNAL DELIVERY



delay



(wall clock)

python_time += interval
c_time += delay - interval

takesometime.py: % of time = 100.00% out of 1.03s.

Line	CPU % Python	CPU % native	Sys %	takesometime.py
1				<code>import time</code>
2				<code>import taketime</code>
3				
4				<code>c_its_per_sec = 250000</code>
5				<code>python_its_per_sec = 1080000</code>
6				
7				<code>import timeit</code>
8				
9				<code>def py_taketime(n):</code>
10				<code> d = 1.01</code>
11				<code> for i in range(n):</code>
12	7%	1%		<code> for j in range(10):</code>
13	18%	3%		<code> d = d * d + d * d</code>
14				<code> return n</code>
15				
16				<code>c_secs = .7</code>
17				<code>python_secs = .3</code>
18				
19				<code>start = time.process_time()</code>
20	1%	70%		<code>p = taketime.run(int(c_secs * c_its_per_sec))</code>
21				<code>q = py_taketime(int(python_secs * python_its_per_sec))</code>
22				<code>end = time.process_time()</code>
23				<code>print(end-start)</code>

≈30%

≈70%


error ≈1%

Other "Irrational exuberance" Languages

Scripting Language	malloc interposition	Monkey patching	Thread enum.	Stack inspection	Opcode disassembly
Perl	✓(1)	✓	<code>threads->list()</code>	<code>Devel::StackTrace</code>	<code>B::Concise</code>
Tcl/Tk	✓(2)	✓			
Python	✓(3)	✓	<code>threading.enumerate()</code>	<code>sys._current_frames()</code>	<code>dis</code>
Lua	✓(4)	✓			
PHP	✓(5)	✓			
R	✓	✓		<code>sys.call</code>	<code>disassemble</code>
Ruby	✓(6)	✓	<code>Thread.list</code>	<code>caller</code>	<code>RubyVM::InstructionSequence</code>

SCALENE


SCRIPTING-LANGUAGE
AWARE
PROFILING



[github.com/emeryberger/](https://github.com/emeryberger/SCALENE) **SCALENE**

SCALENE

SCRIPTING-LANGUAGE
AWARE
PROFILING




[github.com/emeryberger/](https://github.com/emeryberger/SCALENE) **SCALENE**

precise

<u>CPU</u> PYTHON vs. NATIVE + SYS%	<u>MEMORY</u> PYTHON vs. NATIVE	<u>MEMORY</u> USAGE OVER TIME % OF MEM ALLOCATED	<u>COPY</u> VOLUME (MB/s)
--	---------------------------------------	---	---------------------------------

SCALENE

SCRIPTING-LANGUAGE
AWARE
PROFILING




[github.com/emeryberger/](https://github.com/emeryberger/SCALENE) **SCALENE**

precise


fast

<u>CPU</u> PYTHON vs. NATIVE + SYS%	<u>MEMORY</u> PYTHON vs. NATIVE	<u>MEMORY</u> USAGE OVER TIME	<u>COPY</u> VOLUME (MB/s)
		% OF MEM ALLOCATED	



SCALENE

SCRIPTING-LANGUAGE
AWARE
PROFILING




[github.com/emeryberger/](https://github.com/emeryberger/SCALENE) **SCALENE**

precise

fast


<u>CPU</u> PYTHON vs. NATIVE + SYS%	<u>MEMORY</u> PYTHON vs. NATIVE	<u>MEMORY</u> USAGE OVER TIME	<u>COPY</u> VOLUME (MB/s)
		% OF MEM ALLOCATED	



% pip install -U scalene

SCALENE

SCRIPTING-LANGUAGE
AWARE
PROFILING

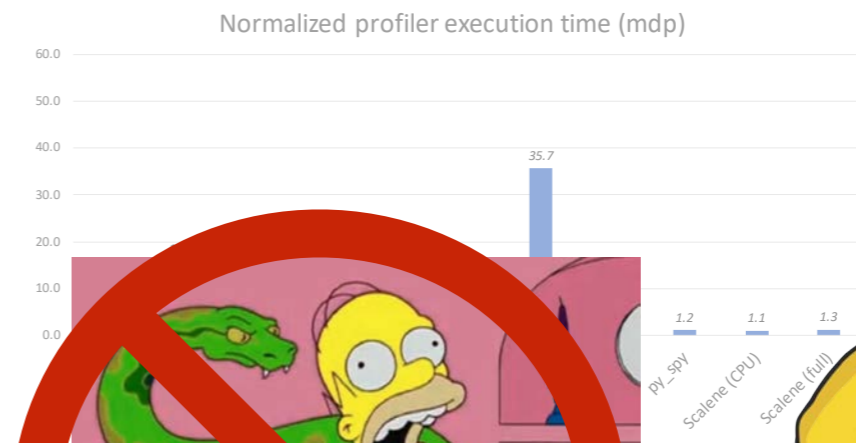


[github.com/emeryberger/](https://github.com/emeryberger/SCALENE) **SCALENE**

precise

fast

<u>CPU</u> PYTHON vs. NATIVE + SYS%	<u>MEMORY</u> PYTHON vs. NATIVE	<u>MEMORY</u> USAGE OVER TIME % OF MEM ALLOCATED	<u>COPY</u> VOLUME (MB/s)
--	---------------------------------------	---	---------------------------------




% pip install -U scalene

SCALENE

SCRIPTING-LANGUAGE
AWARE
PROFILING

github.com/emeryberger/SCALENE



Effective Performance Profiling

COZ

```
% sudo apt install coz-profiler
```

github.com/plasma-umass/coz

