

Delta Pointers: Buffer Overflow Checks Without the Checks

Taddeüs Kroes & Koen Koning
Erik van der Kouwe Herbert Bos
Cristiano Giuffrida

June 19, 2018

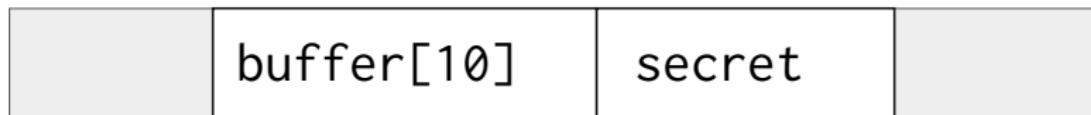


VRIJE
UNIVERSITEIT
AMSTERDAM

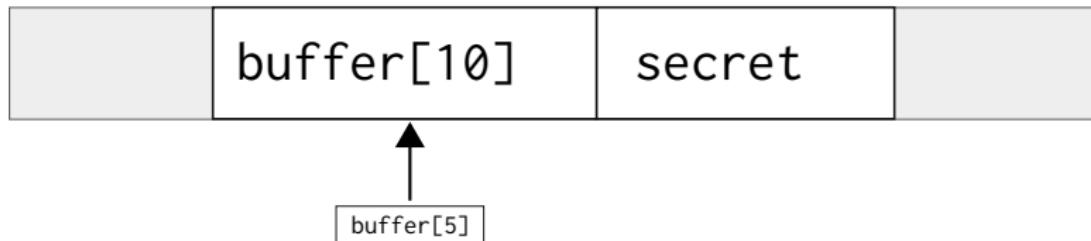


Universiteit
Leiden

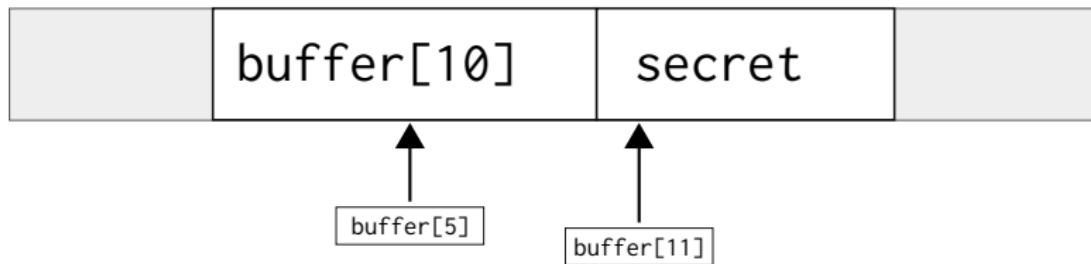
Preview



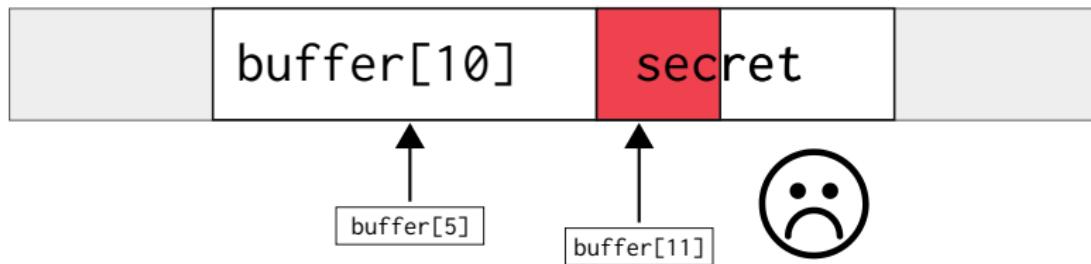
Preview



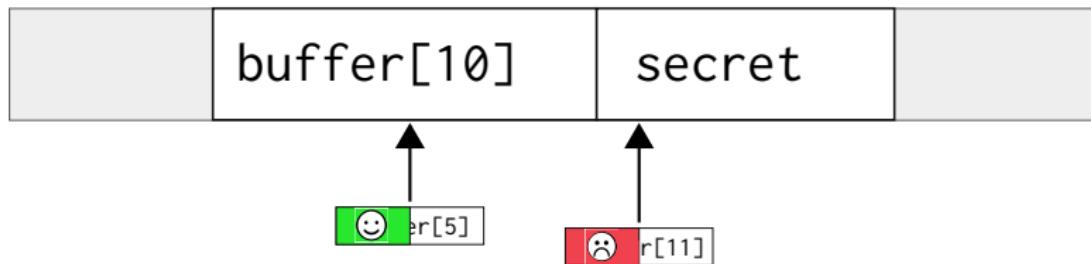
Preview



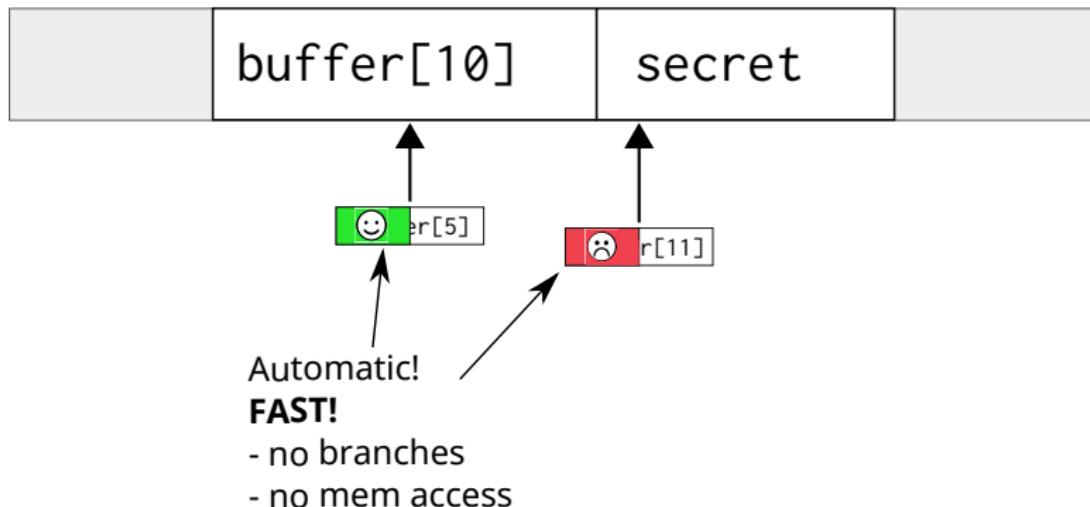
Preview



Preview



Preview



Buffer overflows still very common today



CVE List | CRAN | Board | About | News & Blog



TOTAL CVE Entries: 16433

HOME > CVE > SEARCH RESULTS

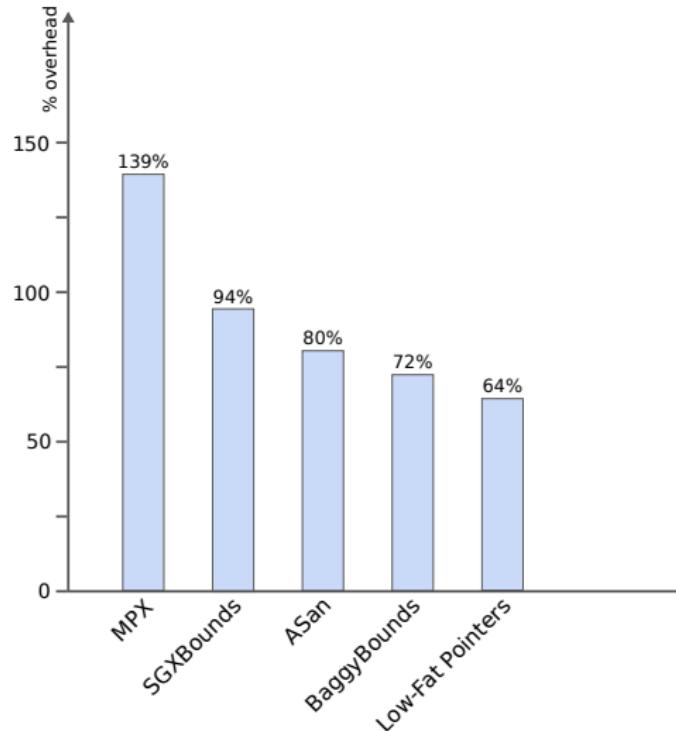
Search Results

There are 9000+ CVE entries that match your search.

Name	Description
CVE-2018-0000	In libTiff 4.0.0, a heap-based buffer overflow occurs in the function LZWDecodeCompat in tl_zinc.c via a crafted TIFF file, as demonstrated by tiffs.
CVE-2018-0001	A Buffer Overflow issue was discovered in Kamailio before 4.4.7, 5.0.x before 5.0.4, and 5.1.x before 5.1.2. A specially crafted REGISTER message with a malformed branch or From tag triggers an off-by-one heap-based buffer overflow in the trc_check_pstr function in modules/trc_mrc_pstr.c.
CVE-2018-0002	The JPTStream::readTEPPart function in JPXStream in spdf-0.0 allows attackers to launch denial of service [heap-based buffer overflow and application crash] or possibly have unspecified other impact via a specific pdf file, as demonstrated by jpstest.
CVE-2018-0003	In PuDefo 0.9.5, there exists a heap-based buffer overflow vulnerability in PuDefo::PuDefenker::GetNextToken() in PuDefenker.cpp, a related issue to CVE-2017-0866. Remote attackers could leverage this vulnerability to cause a denial of service or possibly execute arbitrary code via a crafted pdf file.
CVE-2018-0004	An issue was discovered in CloudBees 1.1.0. An unauthenticated local attacker can connect to the "GoTofile Sync" client application listening on 127.0.0.1 port 8888 and send a malicious payload causing a buffer overflow condition. This will result in code execution, as demonstrated by a TCP reverse shell, or a crash. NOTE: this vulnerability exists because of an incomplete fix for CVE-2018-6682.
CVE-2018-0005	There is a heap-based buffer overflow in the getLength function of utildecapn in libpng 0.4.0 for DOUBLE data. A crafted input will lead to a denial of service attack.
CVE-2018-0006	There is a heap-based buffer overflow in the getLength function of utildecapn in libpng 0.4.0 for INTEGER data. A crafted input will lead to a denial of service attack.
CVE-2018-0007	There is a heap-based buffer overflow in the getLength function of utildecapn in libpng 0.4.0 during a RegisterNumber script. A crafted input will lead to a denial of service attack.
CVE-2018-0008	GRNC through 2.7.1 has a buffer overflow in the gl_media_read_zaa_psnm function in media/tools/psnm.c, a different vulnerability than CVE-2018-1000100.
CVE-2018-0009	An issue was discovered in nplink_m2d_extract in OpenPEPS 2.0.0. The output prefix was not checked for length, which could overflow a buffer, when providing a prefix with 50 or more characters on the command line.
CVE-2018-0010	Stack-based Buffer Overflow in Intel® Tenda AC devices Y1500E 04.12.14_2M allows remote attackers to cause a denial of service or possibly have unspecified other impact.
CVE-2018-0011	There is a heap-based buffer overflow in the pclsCaseRaster function of in_pclsapp in xmpg 4.0.4. A crafted input will lead to a denial of service or possibly have unspecified other impact.
CVE-2018-0012	In Oracle CX ServerSupervision Versions 3.0.0 and prior, parsing malformed project files may cause a stack-based buffer overflow.
CVE-2018-0013	In Oracle CX ServerSupervision Versions 3.0.0 and prior, parsing malformed project files may cause a stack-based buffer overflow.
CVE-2018-0014	In Elixir ECRU versions 2.0.07 and prior, multiple cases where specifically crafted inputs could cause a buffer overflow which, in turn, may allow remote execution of arbitrary code.
CVE-2018-0015	There is a heap-based buffer overflow in the LoadCrx function of in_pcx_rpm in xmpg 4.0.4. A crafted input will lead to a denial of service or possibly unspecified other impact.
CVE-2018-0016	A buffer overflow was found in the MikroTik RouterOS SMB service when processing NetBIOS session request messages. Remote attackers with access to the service can exploit this vulnerability and gain code execution on the system. Authentication takes place, so it is possible for an unauthenticated remote attacker to exploit it. All architectures and all devices running RouterOS before versions 6.41.3%42c27 are vulnerable.
CVE-2018-0017	In usenet2D before 2.3.5, there is a buffer overflow in the usenet_to_araf_parse() function in DriverManager.hdr.c.
CVE-2018-0018	A Buffer Overflow issue was discovered in Apache through 13.0.3, 14.4 through 14.7.5, and 15.1 through 15.2.1, and Certified Asterisks through 13.10.2. When processing a SUBSCRIBE request, the res_gpgsql_pubsub module stored the received data in memory and did not check if the number of headers it processed, despite having a fixed limit of 32. If more than 32 Access headers were present, the code would write outside of its memory and cause a crash.
CVE-2018-0019	The ParseConfigFile function of the config.c file of WinPack 5.1.0 allows a remote attacker to cause a denial of service (global buffer overflow), or possibly trigger a buffer overflow or incorrect memory allocation, via a maliciously crafted CNAME entry.
CVE-2018-0020	An issue was discovered in point2viewer in program:news in Lepotica before 1.7.0. Unsanitized input (hostname) can overflow a buffer, leading potentially to arbitrary code execution or possibly unspecified other impact.
CVE-2018-0021	A buffer overflow vulnerability exist in the web-based GUI of Schneider Electric's Pelco Sans Professional in all firmware versions prior to 2.29.47 which could allow an unauthenticated, remote attacker to execute arbitrary code.
CVE-2018-0022	Lepotica before 1.7.5 does not limit the number of characters in a tsu format argument to hexad or escc, which allows remote attackers to cause a denial of service (stack-based buffer overflow) or possibly have unspecified other impact via a long string, demonstrated by the gethostbyname and gethostid functions.
CVE-2018-0023	Buffer overflow in the decoders function in rtp.c 2.2.0 through 2.4.0(r10) allows remote attackers to execute arbitrary code by leveraging an rtp query and sending a response with a crafted array.
CVE-2018-0024	CONLINE 2.0.0 allows remote attackers to cause a denial of service (buffer overflow) or possibly have unspecified other impact because the cert_md5_pemReadFileB function in cert_pkinfo.c can be called with wrong arguments. Specifically, there is an incorrect integer data type causing a negative third argument in some cases of cert_md5_tlv data with inconsistent length information.
CVE-2018-0025	In CONLINE 2.0.0, the Panner of NENTLV does not verify whether a certain component-length field matches the actual component length, which has a resultant buffer overflow and out-of-bounds memory access.
CVE-2018-0026	In CONLINE 2.0.0, the function cert_md5_to_x509_detailed can cause a buffer overflow, when writing a prefix to the buffer. The maximal size of the prefix is CONL_MAX_PREFIX_SIZE, the buffer has the size CONL_MAX_PREFIX_SIZE. However, when NEN is enabled, additional characters are written to the buffer (e.g., the "NEN" or "RD" tags). Therefore, sending an NEN/RD packet with a prefix of size CONL_MAX_PREFIX_SIZE can cause an overflow of but inside cert_md5_to_x509_detailed.
CVE-2018-0027	An issue was discovered in CloudBees before 1.1.0. An unauthenticated remote attacker can connect to the "GoTofile Sync" client application listening on port 8888 and send a malicious payload causing a buffer overflow condition. This will result in an attacker controlling the programs execution flow and allowing arbitrary code execution.
CVE-2018-0028	An issue was discovered in the basemail function in the SMTP Server in Exim before 4.90.1. By sending a handshaked message, a buffer overflow may happen. This can be used to execute code remotely.
CVE-2018-0029	The userc_exposed_path function in comedit in libkrb5_4.0.0 has a stack-based buffer overflow via a large directory length.
CVE-2018-0030	A stack-based buffer overflow (Remote Code Execution) issue was discovered in Design Science MathType 6.0. This occurs in a function call in which the first argument is a converted off-set value and the second argument is a stack buffer. This is fixed in 6.0.6.
CVE-2018-0031	A buffer overflow vulnerability in the control protocol of Firestone Sync/Breeze Enterprise v10.4.18 allows remote attackers to execute arbitrary code by sending a crafted packet to TCP port 9121.
CVE-2018-0032	A buffer overflow vulnerability in the control protocol of Disk Savvy Enterprise v10.4.18 allows remote attackers to execute arbitrary code by sending a crafted packet to TCP port 9124.
CVE-2018-0033	The pbfreadline function (just like fgets) in libcurl through 0.4.18 is vulnerable to a heap-based buffer overflow, which may allow attackers to cause a denial of service or unspecified other impact via a crafted PDF file.
CVE-2018-0034	Buffer overflows in Hareesa Techno Smartcams



Bounds checking is slow 😞



What is bounds checking?

```
void foo(char *buffer, size_t n) {  
    buffer[n] = 10;  
}
```

What is bounds checking?

```
void foo(char *buffer, size_t n) {  
    buffer[n] = 10;  
}  
  
Attacker-controlled?  
↑
```

What is bounds checking?

```
void foo(char *buffer, size_t n) {  
    if (n >= SIZE(buffer))  
        ERROR("overflow");  
    buffer[n] = 10;  
}  


Automatically  
inserted


```

What is bounds checking?

```
void foo(char *buffer, size_t n) {  
    if (n >= SIZE(buffer))  
        ERROR("overflow");  
    buffer[n] = 10;  
}  
  
Needs  
metadata
```

What is bounds checking?

```
void foo(char *buffer, size_t n) {  
    if (n >= SIZE(buffer))  
        ERROR("overflow");  
    buffer[n] = 10;  
}
```

Branching check }

Needs metadata

What is bounds checking?

```
void foo(char *buffer, size_t n) {  
    if (n >= SIZE(buffer))  
        ERROR("overflow");  
    buffer[n] = 10;  
}
```

Branching check

Overhead!

Needs metadata

```
graph LR; A[Branching check] --> B[if (n >= SIZE(buffer))]; A --> C[Overhead!]; D[Needs metadata] --> E[SIZE(buffer)];
```

What is bounds checking?

```
void foo(char *buffer, size_t n) {  
    if (n >= SIZE(buffer))  
        ERROR("overflow");  
    buffer[n] = 10;  
}
```

Branching check

Overhead!

Needs metadata

Efficient solution:
pointer tagging

The diagram shows annotations on a piece of C code. A red arrow points from the 'if' statement to the text 'Branching check'. Another red arrow points from the assignment statement 'buffer[n] = 10;' to the text 'Overhead!'. A green oval encloses the text 'Needs metadata', with a red arrow pointing to it from the assignment statement. A green arrow points from the oval to the text 'Efficient solution: pointer tagging'.

What is bounds checking?

```
void foo(char *buffer, size_t n) {  
    if (n >= SIZE(buffer))  
        ERROR("overflow");  
    buffer[n] = 10;  
}
```

The code shows a function `foo` that takes a character pointer `buffer` and a size `n`. It checks if `n` is greater than or equal to the size of the buffer. If it is, it calls `ERROR` with the message "overflow". Then it assigns the value 10 to `buffer[n]`.

Annotations:

- A red circle highlights the `if` statement with the text "Branching check" and "Still slow".
- A green circle highlights the assignment `buffer[n] = 10;` with the text "Needs metadata".
- An arrow points from the text "Overhead!" to the `if` statement.
- An arrow points from the text "Efficient solution: pointer tagging" to the green circle.

Our approach: Delta Pointers

- ▶ Use pointer tagging
 - ▶ No memory access for metadata lookup
- ▶ No need for branches
 - ▶ Delegate checks to (off-the-shelf) hardware instead
- ▶ Focus on common case: upper bound on x86_64
 - ▶ Mitigates all CVEs reported by related work

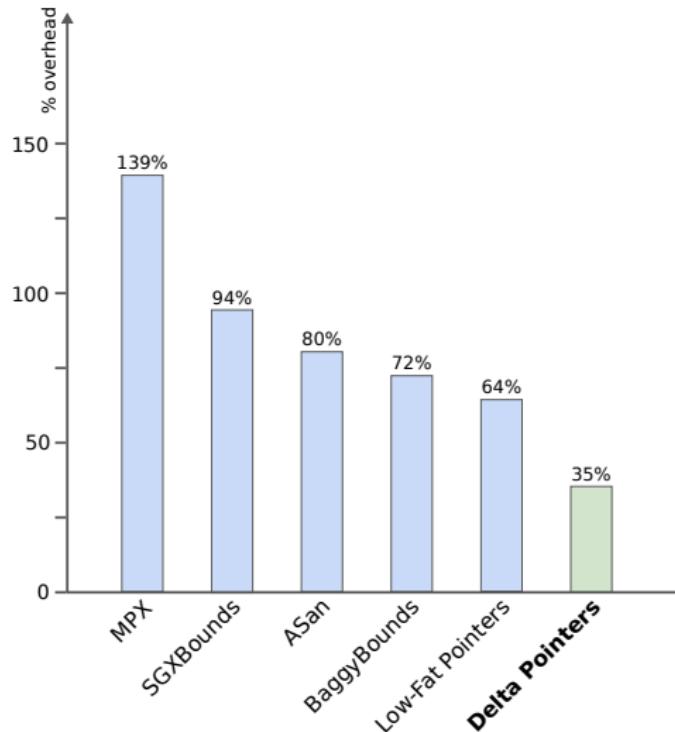
Our approach: Delta Pointers

- ▶ Use pointer tagging
 - ▶ No memory access for metadata lookup
- ▶ No need for branches
 - ▶ Delegate checks to (off-the-shelf) hardware instead
- ▶ Focus on common case: upper bound on x86_64
 - ▶ Mitigates all CVEs reported by related work

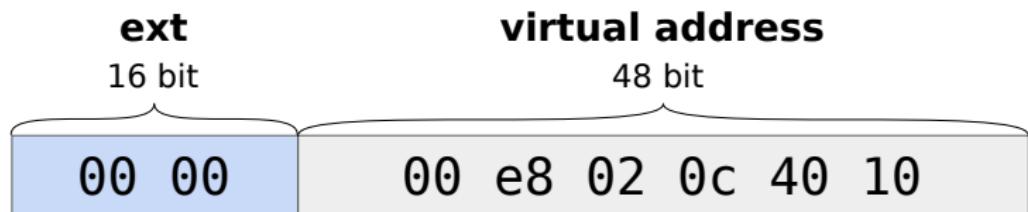
Our approach: Delta Pointers

- ▶ Use pointer tagging
 - ▶ No memory access for metadata lookup
- ▶ No need for branches
 - ▶ Delegate checks to (off-the-shelf) hardware instead
- ▶ Focus on common case: upper bound on x86_64
 - ▶ Mitigates all CVEs reported by related work

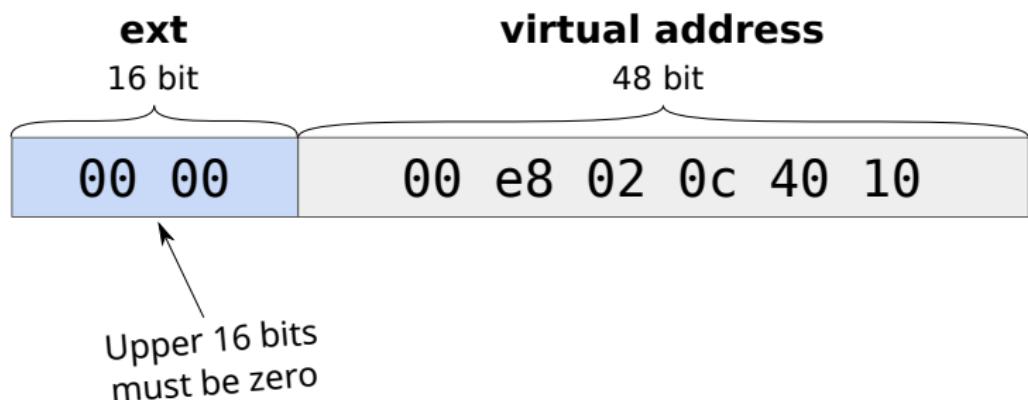
Our approach: Delta Pointers



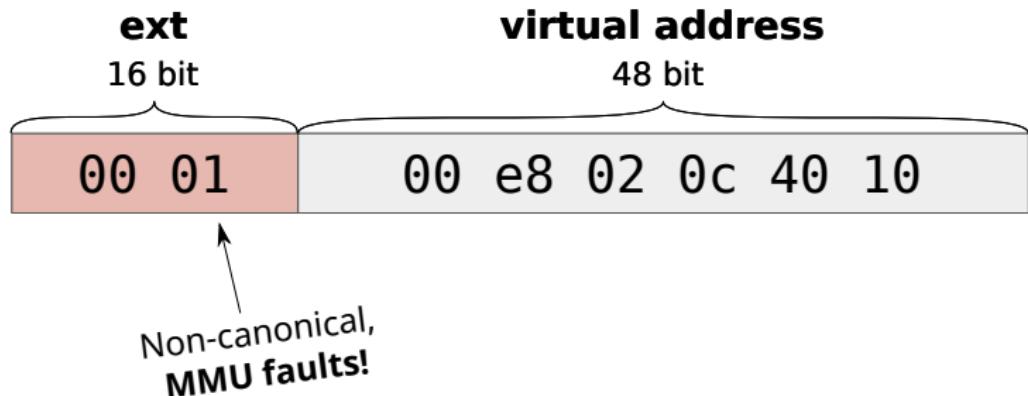
Regular pointers



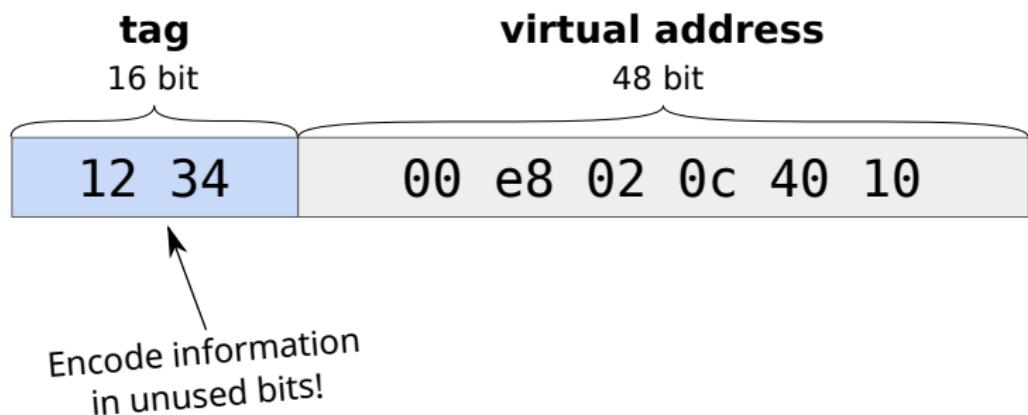
Regular pointers



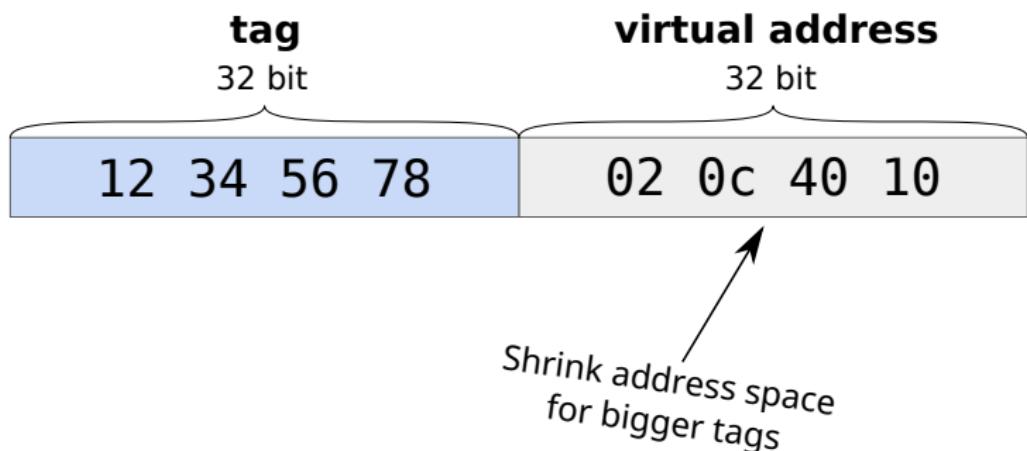
Regular pointers



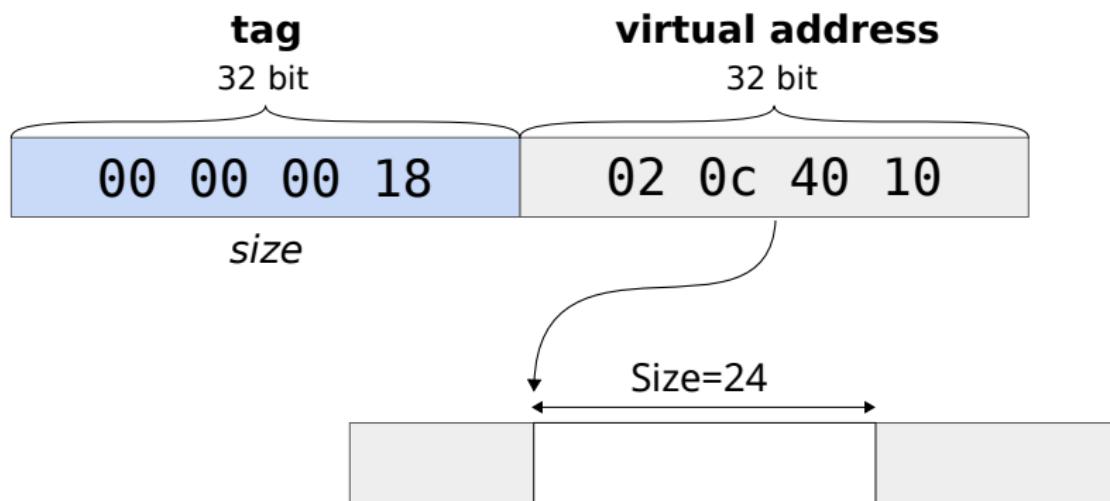
Tagged pointers



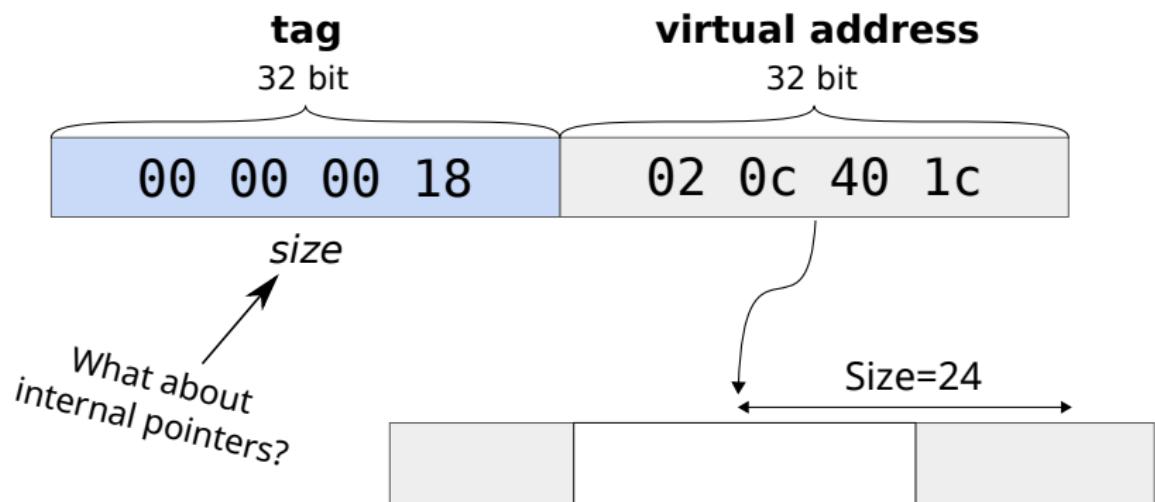
Tagged pointers



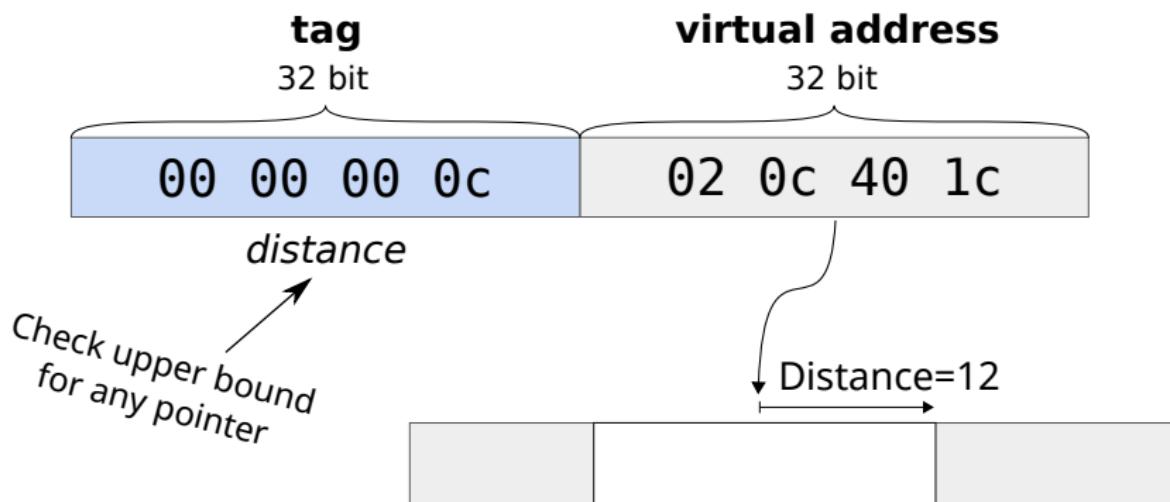
Delta Pointers



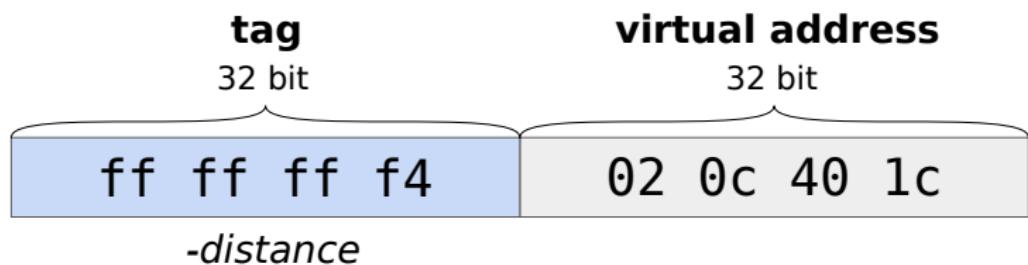
Delta Pointers



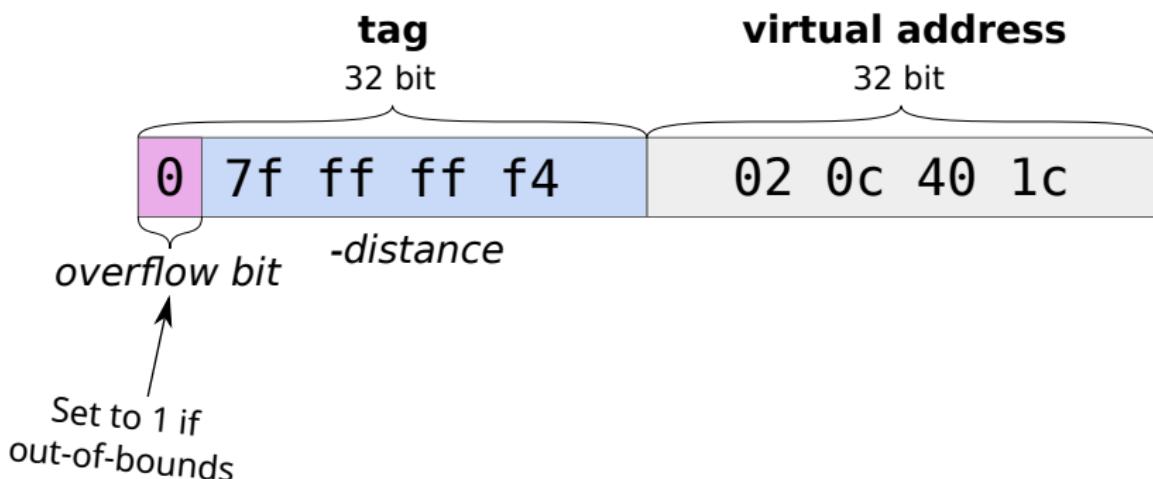
Delta Pointers



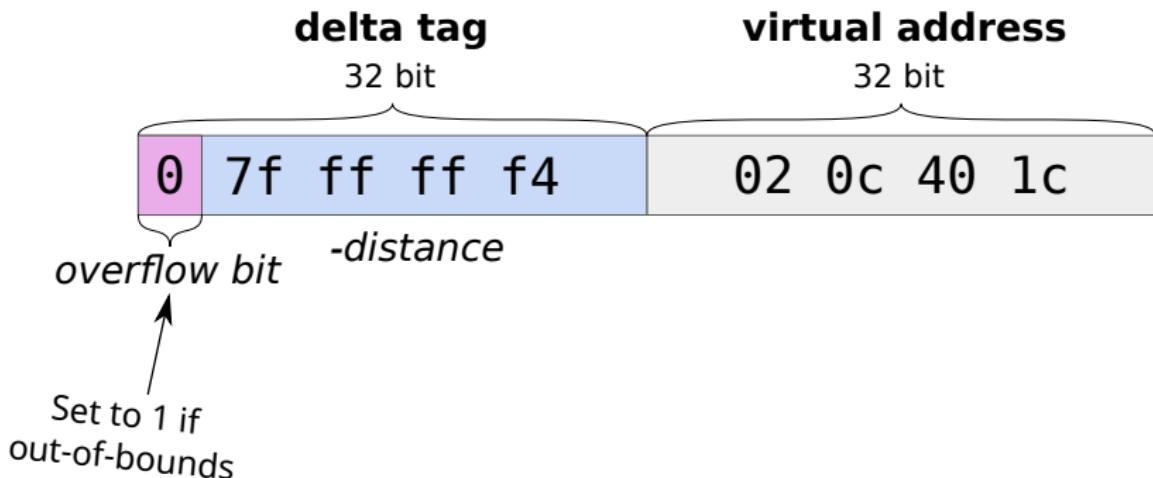
Delta Pointers



Delta Pointers



Delta Pointers



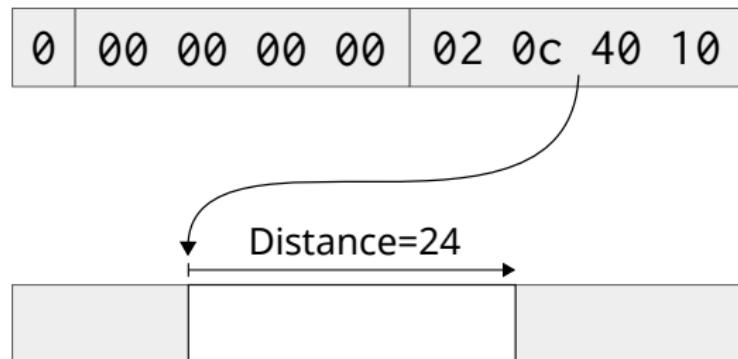
Instrumentation

```
char *p = malloc(24);
```

0	00	00	00	00	02	0c	40	10
---	----	----	----	----	----	----	----	----

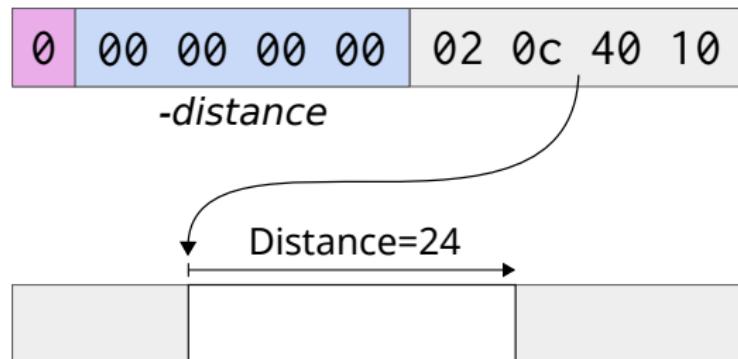
Instrumentation

```
char *p = malloc(24);
```



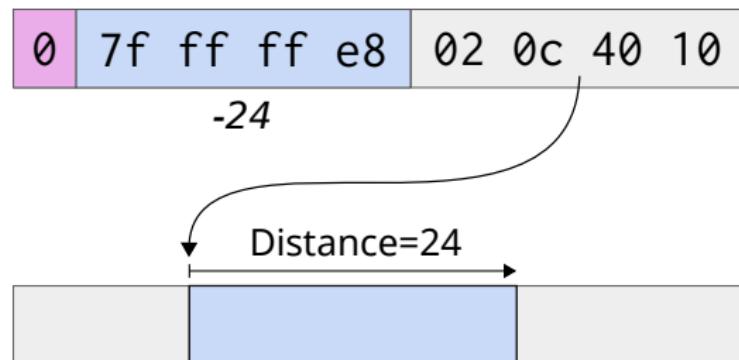
Instrumentation

```
char *p = malloc(24);
```



Instrumentation

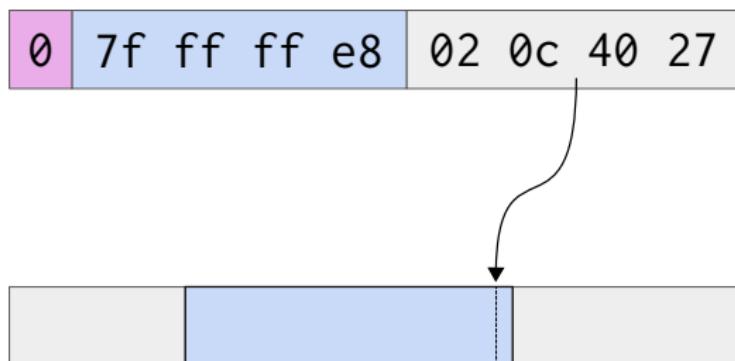
```
char *p = malloc(24) | (-24 << 32);
```



Instrumentation

p += 23;

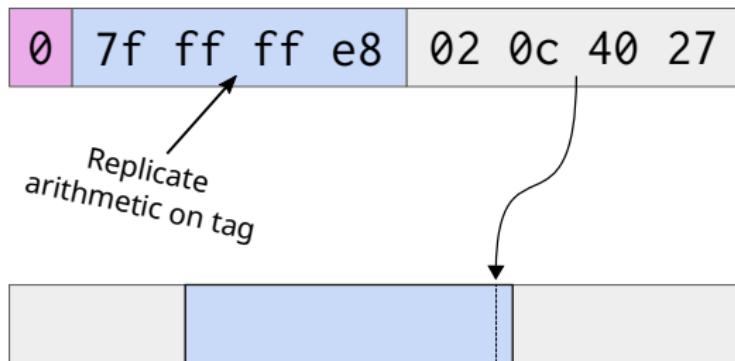
+23



Instrumentation

p += 23;

+23



Instrumentation

```
p += 23 + (23 << 32);
```

+23

+23



Distance=1



Instrumentation

```
p += 1 + (1 << 32);
```

carry

+1

+1



Distance=0,
overflowed!



Instrumentation

```
p += -1 + (-1 << 32);
```



Distance=1,
in-bounds again



Instrumentation

`p += -1 + (-1 << 32); one operation!`



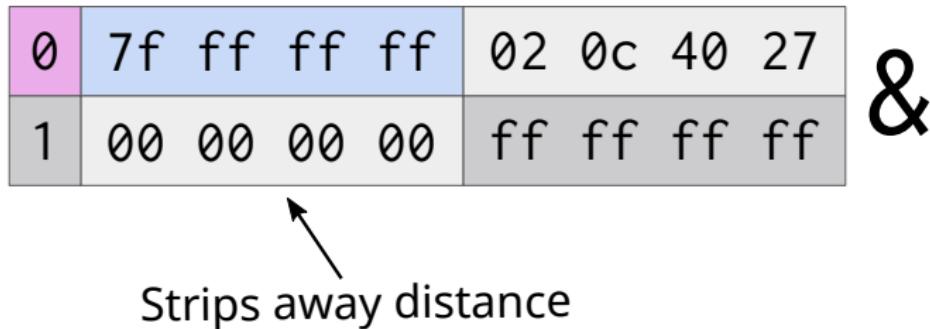
Distance=1,
in-bounds again



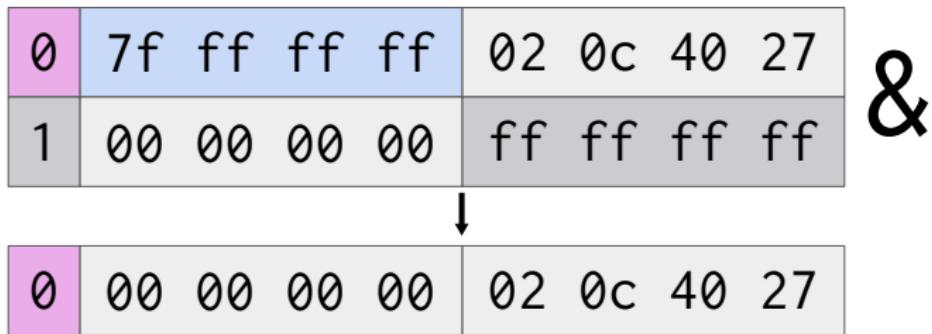
Dereferencing an in-bounds pointer



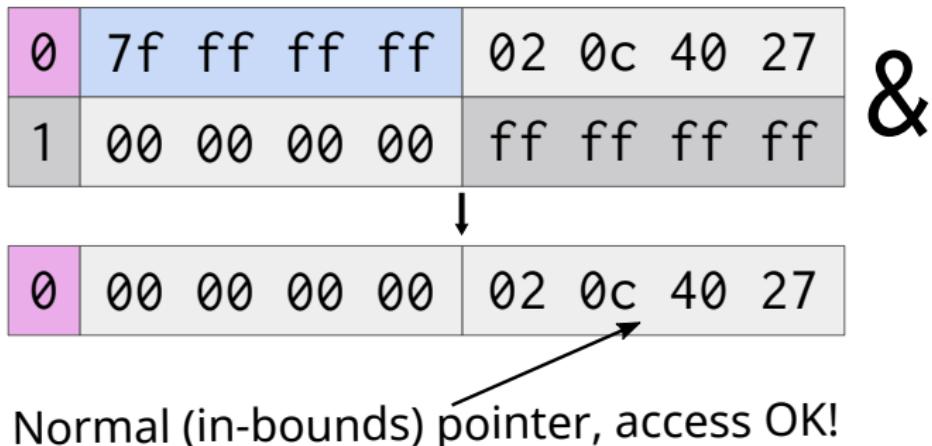
Dereferencing an in-bounds pointer



Dereferencing an in-bounds pointer



Dereferencing an in-bounds pointer



Dereferencing an out-of-bounds pointer

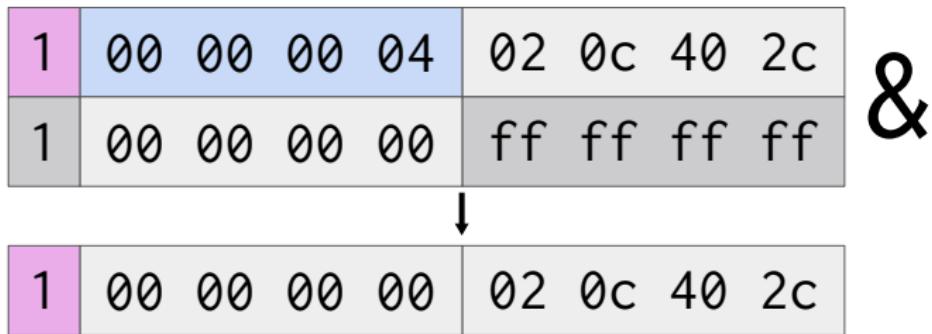


Dereferencing an out-of-bounds pointer

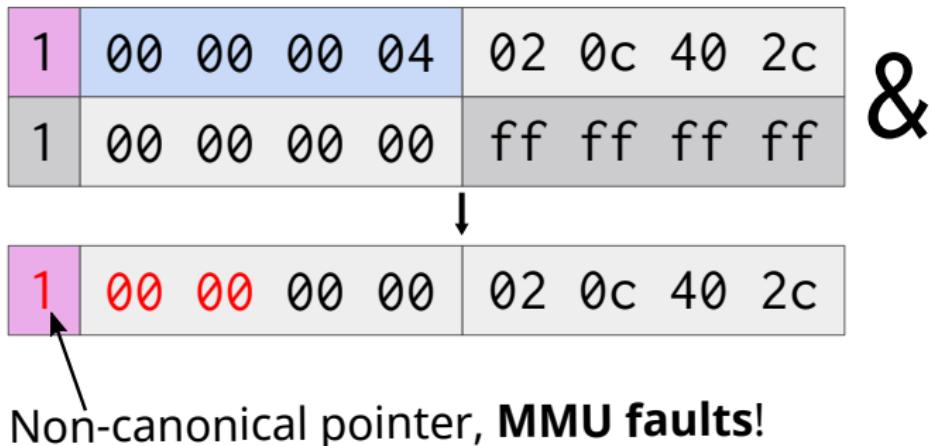
1	00 00 00 04	02 0c 40 2c
1	00 00 00 00	ff ff ff ff

&

Dereferencing an out-of-bounds pointer



Dereferencing an out-of-bounds pointer



Implementation

- ▶ LLVM based prototype for C/C++
- ▶ Stack + heap + globals
- ▶ 32-bit address → 4GB address space
- ▶ 31-bit distance → 2GB allocations
- ▶ Instrument NULL pointer with $distance = -1$
- ▶ Optimizations: omit instrumentation on in-bounds pointers



Pointer tagging breaks things

- ▶ Uninstrumented libraries

```
// strdup(ptr);  
TAG(strdup(MASK(ptr)));
```

- ▶ Non-zero NULL pointer
- ▶ Subtraction, addition, multiplication, vectors, etc.
- ▶ Incomplete type information (e.g., unions)
- ▶ Compiler quirks
- ▶ ... and more
 - ▶ Solved with TBAA + def-use chain analysis
 - ▶ Details in paper



Pointer tagging breaks things

- ▶ Uninstrumented libraries

```
// strdup(ptr);  
TAG(strdup(MASK(ptr)));
```

- ▶ Non-zero NULL pointer

- ▶ Subtraction, addition, multiplication, vectors, etc.

- ▶ Incomplete type information (e.g., unions)

- ▶ Compiler quirks

- ▶ ... and more

- ▶ Solved with TBAA + def-use chain analysis

- ▶ Details in paper



Pointer tagging breaks things

- ▶ Uninstrumented libraries

```
// strdup(ptr);  
TAG(strdup(MASK(ptr)));
```



- ▶ Non-zero NULL pointer
- ▶ Subtraction, addition, multiplication, vectors, etc.
- ▶ Incomplete type information (e.g., unions)
- ▶ Compiler quirks
- ▶ ... and more
 - ▶ Solved with TBAA + def-use chain analysis
 - ▶ Details in paper

Pointer tagging breaks things

- ▶ Uninstrumented libraries

```
// strdup(ptr);  
TAG(strdup(MASK(ptr)));
```

- ▶ Non-zero NULL pointer
- ▶ Subtraction, addition, multiplication, vectors, etc.
- ▶ Incomplete type information (e.g., unions)
- ▶ Compiler quirks
- ▶ ... and more
 - ▶ Solved with TBA + def-use chain analysis
 - ▶ Details in paper



Pointer tagging breaks things

- ▶ Uninstrumented libraries

```
// strdup(ptr);  
TAG(strdup(MASK(ptr)));
```

- ▶ Non-zero NULL pointer
- ▶ Subtraction, addition, multiplication, vectors, etc.
- ▶ Incomplete type information (e.g., unions)
- ▶ Compiler quirks
- ▶ ... and more
 - ▶ Solved with TBA + def-use chain analysis
 - ▶ Details in paper



Pointer tagging breaks things

- ▶ Uninstrumented libraries

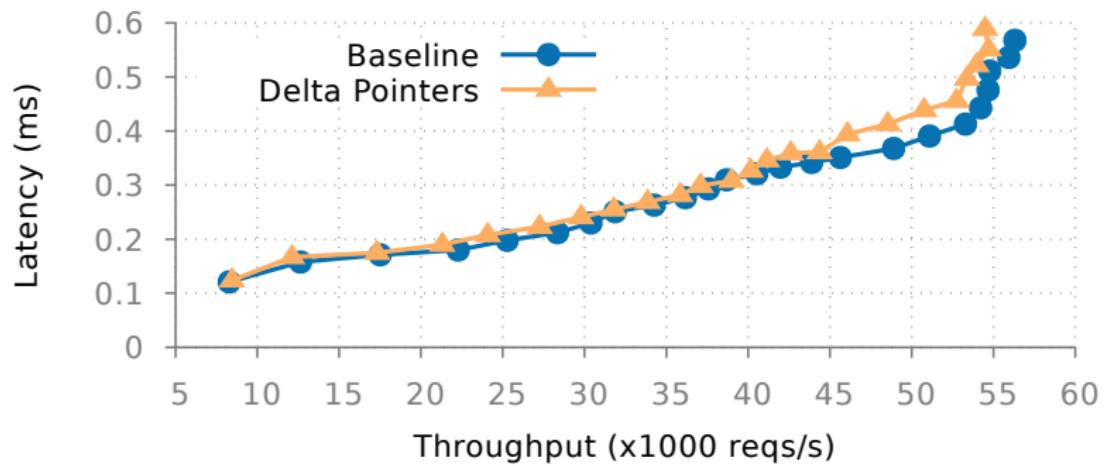
```
// strdup(ptr);  
TAG(strdup(MASK(ptr)));
```

- ▶ Non-zero NULL pointer
- ▶ Subtraction, addition, multiplication, vectors, etc.
- ▶ Incomplete type information (e.g., unions)
- ▶ Compiler quirks
- ▶ ... and more
 - ▶ Solved with TBA + def-use chain analysis
 - ▶ Details in paper



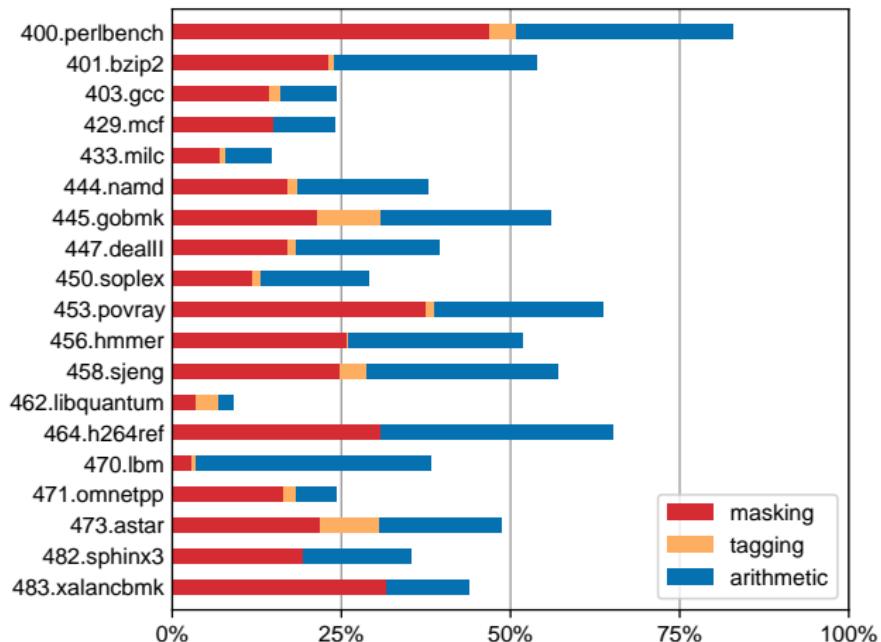
Evaluation

Nginx



3-6% (I/O bound)

SPEC CPU2006 (C/C++)



35% geomean with optimizations

Is that any good?

Is that any good?

Is it better than branches?

Is that any good?

Is it better than branches?

Branching implementation: 48% overhead

Is that any good?

Is it better than branches?

Branching implementation: 48% overhead > 35%!

~~Is that any good?~~

Is it better than branches?

~~Branching implementation: 48% overhead > 35%!~~

Yes

Conclusion

- ▶ Reliable pointer tagging implementation
- ▶ We can check (upper) bounds without checks
- ▶ Faster than existing solutions

<https://github.com/vusec/deltapointers>



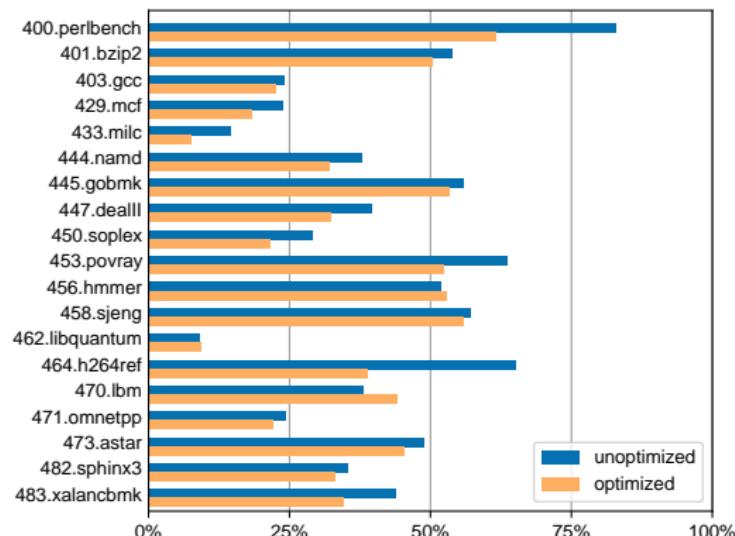
Related work

System	C++	Metadata	Checks	Passing OoB pointers	Non-linear	Runtime	Memory
Softbound	✗	Table	Deref	✓	✓	67%	64%
Baggy Bounds	✗	Layout	Arith	✓ ^a	✓	72%	11%
PArICheck	✗	Shadow	Arith	✓	✓ ^b	96%	18%
LBC	✗	Shadow	Deref	✓	✗	22%	7.7%
ASan	✓	Shadow	Deref	✓	✗	80%	237%
Intel MPX	✓	Table	Deref	✓	✓	139%	90%
LowFat	✓	Layout	Deref	✗	✓	54%	5.2%
SGXBounds	✓	Tag	Deref	✓	✓	89%	0.1%
Delta Pointers	✓	Tag	—	✓	✓	35%	0%

^a Only up to *alloc_size*/2 on 32-bit.

^b Unless wrap-around on 16-bit labels occurs.

Impact of optimization with static analysis



41% \Rightarrow 35%

Statistics

- ▶ 72% of SPEC offsets are dynamic
- ▶ 80% increase in code size with Delta Pointers

Branching implementation

```
void foo(int n) {  
    char *buffer = malloc(24);  
    char *p = buffer + n;  
    *p = 'x';  
}
```

Branching implementation

```
void foo(int n) {  
    char *buffer = malloc(24);  
    buffer |= (buffer + 24) << 32;  
    char *p = buffer + n;  
    *p = 'x';  
}
```

Store end pointer
distance in tag

Branching implementation

```
void foo(int n) {  
    char *buffer = malloc(24);  
    buffer |= (buffer + 24) << 32;  
    char *p = buffer + n;  
    tag = p >> 32; ← Extract tag on  
    p = p & 0xffffffff; ← load/store  
    *p = 'x';  
}
```

Branching implementation

```
void foo(int n) {  
    char *buffer = malloc(24);  
    buffer |= (buffer + 24) << 32;  
    char *p = buffer + n;  
    tag = p >> 32;  
    p = p & 0xffffffff;  
    Branching  
    check → if (p >= tag)  
            ERROR("overflow");  
    *p = 'x';  
}
```

Some pointer tagging challenges

- ▶ Some operations need masking to preserve semantics

```
char a[10];
//size_t len = &a[10] - &a[0];
size_t len = MASK(&a[10]) - MASK(&a[0]);
```

- ▶ Pointers that look like integers

```
union {
    char *buf;
    uint64_t foo;
} field;

field.buf += 42; // should instrument
field.foo += 42; // should NOT
```