

```

1541 1 class T {
1542 2     long a = 0;
1543 3     int b() {
1544 4         int d = 2, e = 10;
1545 5         for (int m = 7; m < 149; ++m)
1546 6             for (int c = 1; c < 4; c++)
1547 7                 for (int n = 1; n < 2; ++n)
1548 8                     d += e;
1549 9         return d;
1550 10    }
1551 11    void f() {
1552 12        for (int w = 0; w < 6361; w++)
1553 13            a = b();
1554 14    }
1555 15    public static void main(String[] g) {
1556 16        T t = new T();
1557 17        for (int h = 1; h < 212; h++)
1558 18            t.f();
1559 19        System.out.println(t.a);
1560 20    }
1561 21 }

```

**Figure 4. Issue-\*\*\*\*\*:** Mis-compilation. The highlighted neutral loop is inserted by Artemis. Code shown in this example is a cleaned-up version from a very large test program.

## A More Examples

Artemis is fruitful in finding diverse bugs such as segmentation faults (SIGSEGV), fatal arithmetic error (SIGFPE), emergency abort (SIGABRT), assertion failures, mis-compilations, and performance issues. We discuss a small selection to highlight the diversity. All tests shown in this section originate from JavaFuzzer [19] and are mutated by Artemis; they are reduced from much larger mutants with a combination of Perses, C-Reduce, and further manual cleanup if needed.

### A.1 Illustrative OpenJ9 Example

Figure 4 triggers a mis-compilation in OpenJ9. Artemis found the bug in OpenJ9 0.32 (revision 3d06b2f9c, based on OpenJDK 1.8.0\_342). However, it can be reproduced at least as far back as 0.24 with JDK 11 and immediately marked as blocker, the most severe, release-blocking type of bugs.

The seed program assigns `T.a` by a value returned by `T.b()` repeatedly (Line 13). Because `T.b()` is pure, it should always return a fixed integer 4262 and the program output should be 4262 as well. For the seed program, the method counter of `T.b()` reaches the warm-level compilation threshold and `T.b()` is JIT-compiled at warm level; all other methods are barely interpreted until the program exits. Artemis alters the JIT compilation by inserting the highlighted loop into Line 12, which leads to an additional JIT compilation of `T.f()` at warm level and two additional OSR compilations of the inserted loop at veryHot and scorching levels, respectively. Considering `T.b()` constantly returns 4262, the inserted loop is neutral and the insertion does not affect the seed's semantics. Therefore, the mutant should output 4262 as the seed. Yet, OpenJ9's JIT compiler mis-compiles the mutant and outputs 1422.

```

1596 1 class T {
1597 2     public static void main(String[] g) {
1598 3         for (int i=0; i<10; i+=1)
1599 4             for (int j=1; j<197; j+=1)
1600 5                 for (int k=-4910; k<314; k+=1)
1601 6                     try {
1602 7                         boolean b = false;
1603 8                         byte[] a = new byte[1 << 14];
1604 9                         try (ByteArrayOutputStream o =
1605 10                             new ByteArrayOutputStream();
1606 11                             ZipOutputStream z =
1607 12                                 new ZipOutputStream(o)) {
1608 13                             final byte[] x = { 'x' };
1609 14                             z.putNextEntry(new ZipEntry("a.gz"));
1610 15                             GZIPOutputStream g = new GZIPOutputStream(z);
1611 16                             g.write(x); g.finish();
1612 17                             if (b) z.write(x);
1613 18                             z.closeEntry();
1614 19                             z.putNextEntry(new ZipEntry("b.gz"));
1615 20                             GZIPOutputStream k = new GZIPOutputStream(z);
1616 21                             k.write(x); k.finish();
1617 22                             z.closeEntry(); z.flush();
1618 23                             a = o.toByteArray();
1619 24                         }
1620 25                     } catch (Throwable x) {}
1621 26     }
1622 27 }

```

**Figure 5. JDK-\*\*\*\*\*:** Performance Issue. When executing the C2 compiled code, the HotSpot process running this test is directly killed on Ubuntu while becomes much slower on Windows. Code shown in this example is a cleaned-up version from a very large test program.

The root cause is that the Expressions Simplification pass is buggy when considering whether an expression is an invariant which should be hoisted out of its residing loop. This causes the JIT compiler to calculate an incorrect number of iterations (142 instead of 426) when hoisting the expression at Line 8 out of the loop at Line 5. To fix this, the developers updated the invariant-check policy.

It should be noted that this bug manifests merely when `T.b()` is JIT-compiled and `T.f()` is hot enough such that the inserted loop is OSR-compiled at scorching level. That said, JIT-compiling `T.b()` and scorchingly JIT-compiling `T.f()` of the seed program<sup>7</sup> cannot trigger the bug as it does not involve OSR. Although it is theoretically possible for a test generator to generate such a bug-triggering program without Artemis, the time budget is generally unpredictable. As a comparison, Artemis can trigger this bug within 8 mutations of the seed program.

### A.2 More Examples

**Figure 6a.** When HotSpot C1 compiles the given code, it tries to inline the method `invokeExact()` but requires the receiver to be non-null (Line 14). HotSpot developers forgot the null-check, and it thus triggers an assertion failure.

<sup>7</sup>-Xjit:{T.b()I}(count=0),{T.f()V}(count=0,optLevel=scorching)

**Figure 6b.** The DLT (Dynamic Loop Transfer which facilitates OSR) optimization pass in OpenJ9 fails to preserve the type information of the byte array `e` (Line 7), leading to an unexpected byte-bit truncation.

**Figure 6c.** OpenJ9 crashes with a segmentation fault when the method `T.e()` is compiled at scorching level. This is because OpenJ9's JIT compiler accesses a removed or invalid block when traversing predecessor blocks to compute and extend the live ranges of some variables.

**Figure 6d.** When compiling the loop in `T.f()`, the HotSpot C2 generates incorrect code, making HotSpot crash with a fatal arithmetic error when running the compiled code.

**Figure 6e.** OpenJ9 fails to vectorize the loop in `T.p()` and as a result, crashes as a segmentation fault because of a null pointer dereference when checking for loop independence.

**Figure 6f.** When an array access is out-of-bound in the compiled code, ART should de-optimize along a bounds-check slow path where the exception should be caught if there is an exception handler. However, ART's bounds-check

slow path generator fails to compute the correct array index and length, which causes the exception to be caught at an incorrect index.

**Figure 5.** This presents the only performance issue found by Artemis. When running the code OSR-compiled by HotSpot C2, the native memory keeps increasing unexpectedly. This finally drains the machine of its memory and the process running this test is thereby killed by the underlying Ubuntu system (16GiB RAM). In contrast, despite not being killed on Windows (16GiB RAM), the process becomes far slower even than interpreting the bytecode. This issue was confirmed as a potential memory leak bug immediately after we reported, but marked as "Won't Fix" after several weeks because the developers hypothesized that "the app is simply allocating memory faster than the GC can reclaim it." However, we believe this issue should be a real bug that affects end users because (1) GC should be opaque to the end users, and (2) it should not be necessary for end users to be concerned with how fast the GC is when developing Java code.

```

1761 1 class T {
1762 2   boolean b;
1763 3
1764 4   void a() {
1765 5     int c, v;
1766 6     double d = 2.61331;
1767 7     for (int z=830; z>51; --z)
1768 8       b = b;
1769 9     v = 1;
1770 10    while (++v < 908)
1771 11      for (c=-3230; c<9840; c+=2)
1772 12        try {
1773 13          MethodHandle m = null;
1774 14          m.invokeExact();
1775 15        } catch (Throwable t) {
1776 16        } finally {}
1777 17        System.out.println(d);
1778 18      }
1779 19    }
1780 20    public static void main(...) {
1781 21      T t = new T();
1782 22      for (;;)
1783 23        t.a();
1784 24    }
1785 25  }

```

(a) JDK-\*\*\*\*: Assertion Failure

```

1786 1 class T {
1787 2   long c; int x;
1788 3   void f() {
1789 4     int i, p, q = 42252;
1790 5     int j, g, k = 5;
1791 6     double m;
1792 7     long[] l = new long[256];
1793 8     for (int o=2; o<87;) {
1794 9       boolean z = false;
1795 10      for (j=737; j<16822; j+=3)
1796 11        if (!z) {
1797 12          z = true;
1798 13          for (p=1; p<6; ++p) {
1799 14            for (m=1; m<2; m++) {
1800 15              l[(int) m] -= 16;
1801 16              c >>= o;
1802 17            }
1803 18            for (g=1; 2>g; g+=3) {
1804 19              l[g] -= k; i = (int) c;
1805 20              try { k = q / i; }
1806 21              catch (ArithExcp w) {}
1807 22            }
1808 23            switch (o) {
1809 24              case 73: x = p;
1810 25            }
1811 26          }
1812 27        }
1813 28      }
1814 29    }
1815 30    void h() { f(); }
1816 31    public static void main(...) {
1817 32      T t = new T(); t.h();
1818 33    }
1819 34  }

```

(d) JDK-\*\*\*\*: Fatal Arithmetic Error

```

1820 1 class T {
1821 2   void l() {
1822 3     int m, d;
1823 4     byte[] e = new byte[6];
1824 5     for (int j=2; j<222; ++j) {
1825 6       for (m=d=1; d<4; d+=2)
1826 7         e[m] -= d;
1827 8       for (int z=5; z<948; z++) {
1828 9         boolean[] x = new boolean[576];
1829 10      }
1830 11    }
1831 12    System.out.println(f(e));
1832 13  }
1833 14  long f(byte[] a) {
1834 15    long k = 0; int i = 0;
1835 16    while (i < a.length) {
1836 17      k += a[i]; i++;
1837 18    }
1838 19    return k;
1839 20  }
1840 21  public static void main(...) {
1841 22    T t = new T(); t.l();
1842 23  }
1843 24  }
1844 25  }

```

(b) Issue-\*\*\*\*: Mis-compilation

```

1845 1 class T {
1846 2   long i;
1847 3
1848 4   void p(long l) {
1849 5     int f = 0;
1850 6     int g[] = new int[0];
1851 7     long[] h = new long[0];
1852 8     for (int j=7; j<84;)
1853 9       try {
1854 10        long[] a = {1};
1855 11        for (int z=1; z<=f; z++)
1856 12          l += a[z - 1];
1857 13      } catch (Throwable r) {
1858 14      } finally {
1859 15        l = 7;
1860 16      }
1861 17      i = 1 + j(g) + j(h);
1862 18    }
1863 19  }
1864 20  void k() { p(0l); }
1865 21  long j(long[] a) { return 0; }
1866 22  long j(int[] a) { return 0; }
1867 23
1868 24  public static void main(...) {
1869 25    T t = new T();
1870 26    t.k();
1871 27  }
1872 28  }
1873 29  }
1874 30  }
1875 31  }
1876 32  }
1877 33  }
1878 34  }

```

(e) Issue-\*\*\*\*: Segmentation Fault

```

1879 1 class T {
1880 2   void e() {
1881 3     for (int i=-4605; i<2; i++)
1882 4       try {
1883 5         double[] d = new double[1];
1884 6         d[0] = 0;
1885 7       } catch (Throwable x) {
1886 8       }
1887 9     }
1888 10  }
1889 11  public static void main(...) {
1890 12    T s = new T();
1891 13    for (;;)
1892 14      s.e();
1893 15  }
1894 16  }
1895 17  }
1896 18  }
1897 19  }
1898 20  }
1899 21  }
1900 22  }
1901 23  }
1902 24  }
1903 25  }

```

(c) Issue-\*\*\*\*: Segmentation Fault

```

1904 1 class T {
1905 2   int r;
1906 3
1907 4   void f() {
1908 5     int i, j, o = 5788, t = 127;
1909 6     byte[] a = new byte[400];
1910 7     for (i=14; 297>i; ++i)
1911 8       for (j=151430; j<235417; j+=2);
1912 9     try {
1913 10      for (int d=4; 179>d; ++d) {
1914 11        o *= o;
1915 12        for (int k=1; k<58; k++)
1916 13          for (int z=k; 1+400>z;
1917 14            z++) {
1918 15            a[z] -= t;
1919 16            o += z;
1920 17            switch (d % 5) {
1921 18              case 107: d >>= r;
1922 19            }
1923 20          }
1924 21        } catch (AI00BExcp x) {}
1925 22      System.out.println(i + "," + o);
1926 23    }
1927 24  }
1928 25  }
1929 26  public static void main(...) {
1930 27    T t = new T(); t.f();
1931 28  }
1932 29  }
1933 30  }
1934 31  }
1935 32  }
1936 33  }
1937 34  }

```

(f) Issue-\*\*\*\*: Mis-compilation

**Figure 6.** A small selection of example tests finding various JVM bugs. For simplicity, we omit the required imports and the parameter of `T.main(String[] args)` method. Acronym `ArithExcp` represents `ArithmeticException`. Acronym `AI00BExcp` stands for `ArrayIndexOutOfBoundsException`. Code shown in these examples are cleaned-up versions from very large test programs generated with a combination of `JavaFuzzer` and `Artemis`.