

# Java NIO core knowledge summary

👤 [Guide](#) 📄 Java 📌 Java IO, Java Basics 📖 About 3899 words ⌚ About 13 minutes

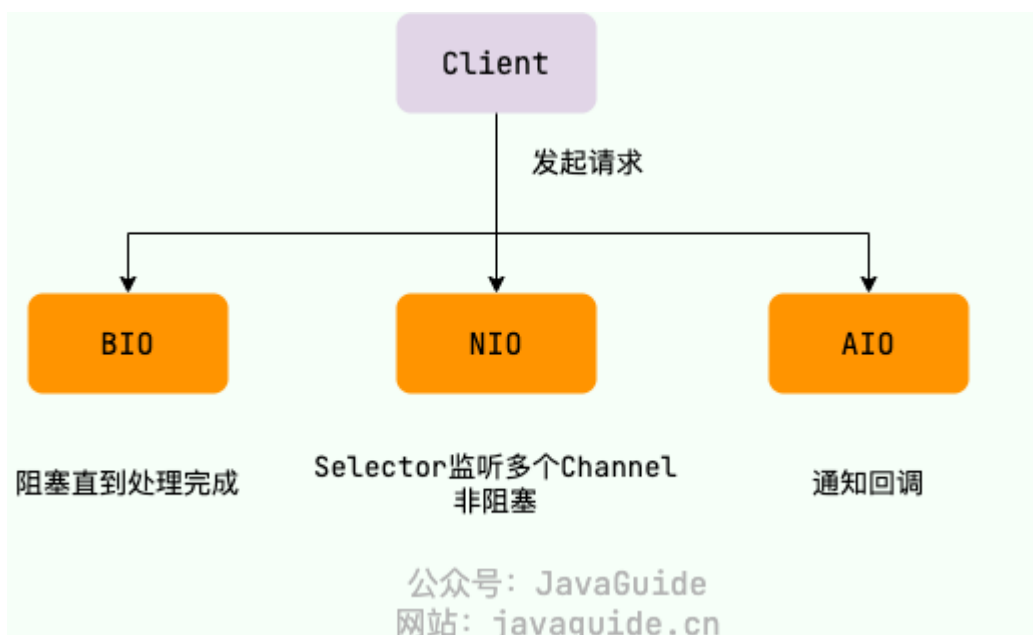
Before learning NIO, you need to understand the basic theoretical knowledge of the computer I/O model. If you don't understand it yet, you can refer to this article I wrote: [Detailed Explanation of the Java IO Model](#) .

## About NIO

In the traditional Java I/O model (BIO), I/O operations are performed in a blocking manner. That is, when a thread performs an I/O operation, it is blocked until the operation completes. This blocking model can lead to performance bottlenecks when handling multiple concurrent connections because a thread must be created for each connection, and thread creation and switching are both expensive.

To address this issue, Java 1.4 introduced a new I/O model— **NIO** (New IO, also known as Non-blocking IO). NIO addresses the shortcomings of synchronous blocking I/O by providing non-blocking, buffered, channel-based I/O within standard Java code. It can handle multiple connections using a small number of threads, significantly improving I/O efficiency and concurrency.

The figure below is a simple comparison of BIO, NIO and AIO processing client requests (for an introduction to AIO, you can read this article I wrote: [Detailed Explanation of Java IO Model](#) , which is not the focus, just understand it).



⚠️ **Note:** Using NIO doesn't necessarily guarantee high performance. Its performance advantage is primarily in high-concurrency and high-latency network environments. When there are fewer connections, lower concurrency, or faster network transmission speeds, NIO's performance may not necessarily outperform traditional BIO.

## NIO core components

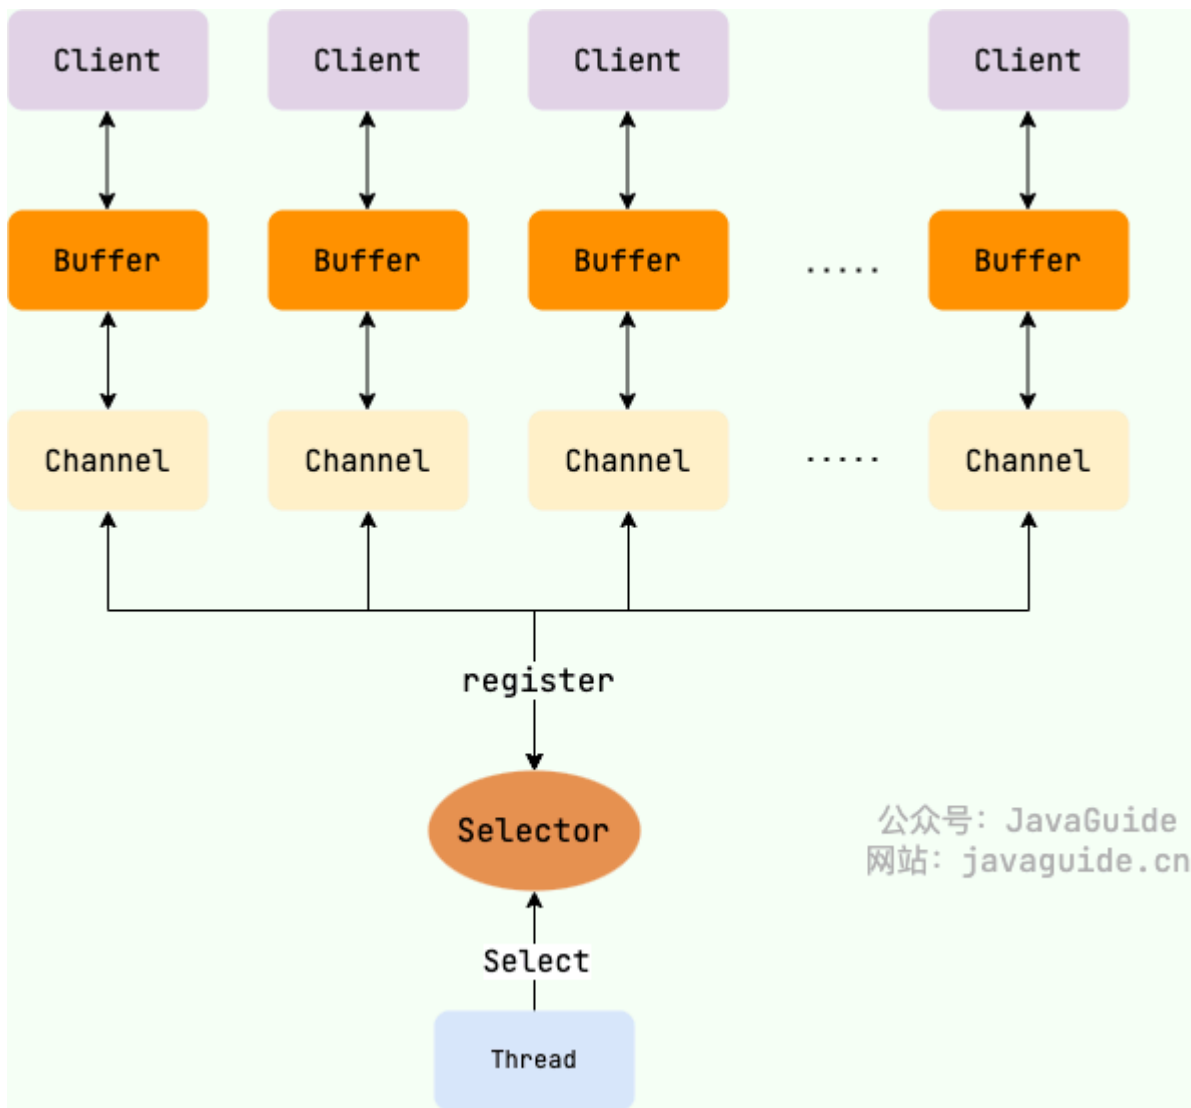
---

NIO mainly includes the following three core components:

- **Buffer** : NIO reads and writes data through the buffer. During a read operation, the data in the channel is filled into the buffer, and during a write operation, the data in the buffer is written into the channel.
- **Channel** : A channel is a bidirectional, readable and writable data transmission channel. NIO uses channels to implement data input and output. A channel is an abstract concept that can represent a connection between files, sockets, or other data sources.
- **Selector** : Allows a single thread to handle multiple channels, based on an event-driven I/O multiplexing model. All channels can be registered with a Selector, which then assigns threads to handle events.

The relationship between the three is shown in the figure below (it doesn't matter if you don't understand it for now, it will be explained in detail later):





The following is a detailed introduction to these three components.

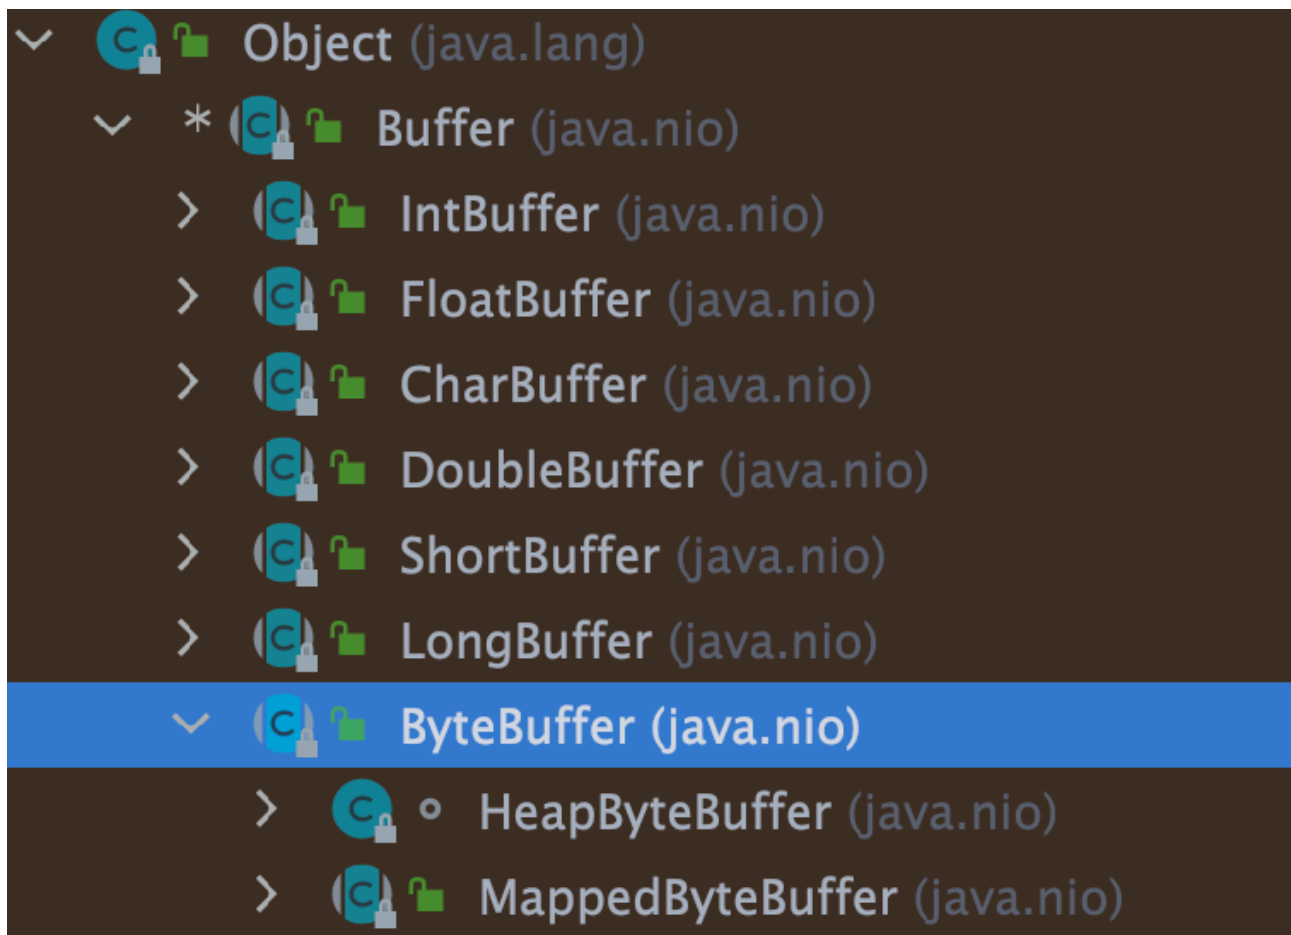
## Buffer

In traditional BIO, data reading and writing are stream-oriented, divided into byte stream and character stream.

In Java 1.4's NIO library, all data is handled using buffers. This is a key difference between the new library and the previous BIO library, somewhat similar to the buffered streams in BIO. When reading data, NIO reads it directly into the buffer. When writing data, it writes it directly into the buffer. When using NIO, both reading and writing data are performed through the buffer.

Buffer The subclasses of are shown in the figure below. Among them, the most commonly used is `ByteBuffer`, it can be used to store and operate byte data.





You can think of Buffer as an array, IntBuffer where, FloatBuffer, etc. correspond to int[], float[], char[], etc. CharBuffer respectively.

To understand the buffer more clearly, let's take a quick look at Buffer the four member variables defined in the class:

```

1  public abstract class Buffer {
2      // Invariants: mark <= position <= limit <= capacity
3      private int mark = -1;
4      private int position = 0;
5      private int limit;
6      private int capacity;
7  }

```

java

The specific meanings of these four member variables are as follows:

1. Capacity ( capacity ): Buffer The maximum amount of data that can be stored, Buffer set at creation time and cannot be changed;
2. Limit ( limit ): Buffer The boundary of the data that can be read/written. In write mode, limit it represents the maximum amount of data that can be written, generally

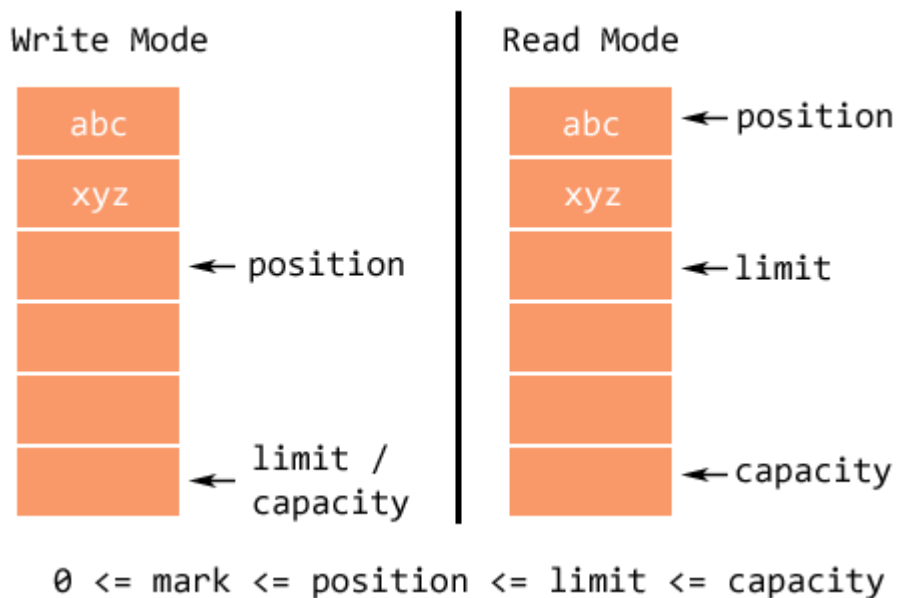


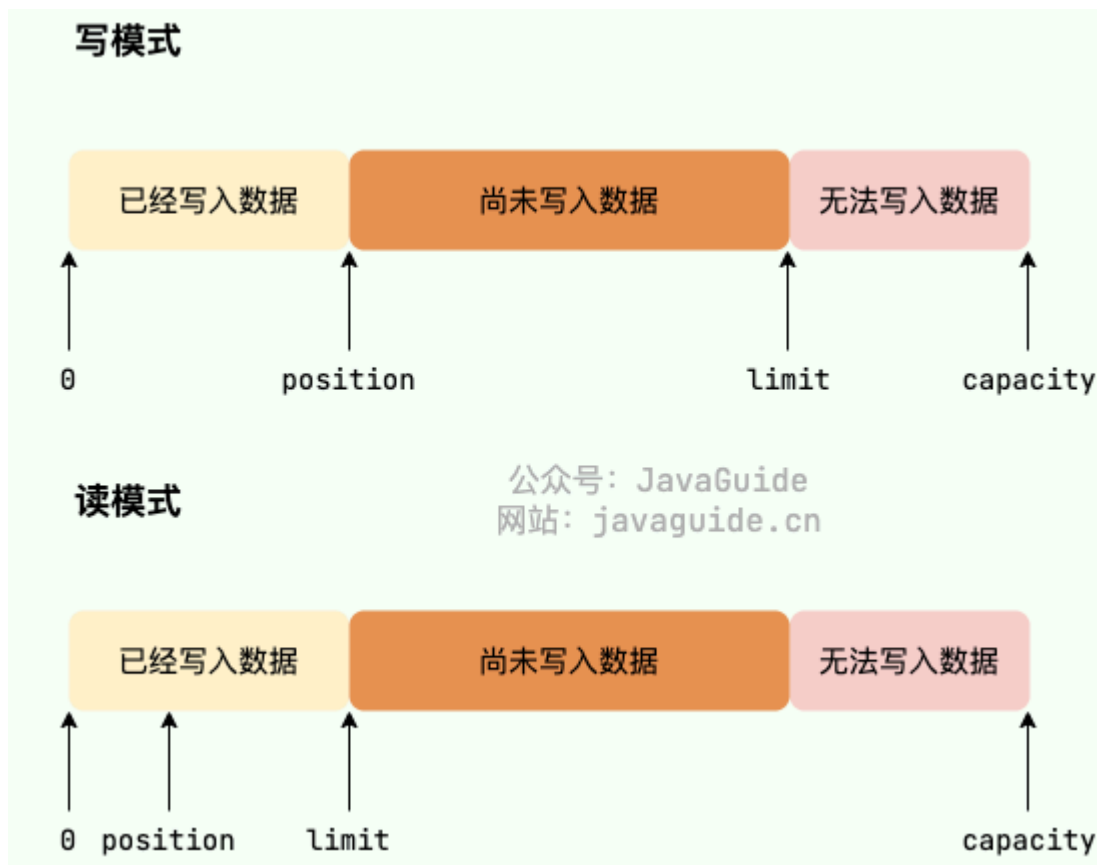
equal to `capacity` (can `limit(int newLimit)` be set through methods); in read mode, `limit` it is equal to the actual size of the data written to the buffer.

3. `Position ( position )`: The position (index) of the next data that can be read or written. When switching from write mode to read mode (`flip()`), `position` it will be reset to zero so that reading and writing can start from the beginning.
4. `Marker ( mark )`: Buffer allows the position to be directly positioned at the marker, which is an optional attribute;

Moreover, the above variables satisfy the following relationship:  **$0 \leq \text{mark} \leq \text{position} \leq \text{limit} \leq \text{capacity}$** .

Buffers have two modes: read mode and write mode, used to read and write data to the buffer, respectively. When a buffer is created, it is in write mode by default `flip()`. You can switch to read mode by calling `Buffer()`. To switch back to write mode, you can call the `Buffer()` or `Buffer clear() () compact()` methods.





Buffer Objects cannot new be created by calling constructors and can only be instantiated through static methods `Buffer` .

Here `ByteBuffer` we take as an example:

```

1  //
2  public static ByteBuffer allocate(int capacity);
3  //
4  public static ByteBuffer allocateDirect(int capacity);

```

java

The two core methods of Buffer are:

1. `get` : Read data from the buffer
2. `put` : Write data to the buffer

In addition to the above two methods, other important methods are:

- `flip` : Switches the buffer from write mode to read mode, which `limit` sets the value of to the current `position` value of and `position` sets the value of to 0.
- `clear` : Clear the buffer, switch the buffer from read mode to write mode, and `position` set the value of to 0 and `limit` the value of to `capacity` the value of.
- ...

## The process of data change in Buffer:

```

1  import java.nio.*;
2
3  public class CharBufferDemo {
4      public static void main(String[] args) {
5          //          8 CharBuffer
6          CharBuffer buffer = CharBuffer.allocate(8);
7          System.out.println("          ");
8          printState(buffer);
9
10         //  buffer  3
11         buffer.put('a').put('b').put('c');
12         System.out.println("    3          ");
13         printState(buffer);
14
15         //      flip()          buffer          position
16         0, limit    3
17         buffer.flip();
18         System.out.println("    flip()          ");
19         printState(buffer);
20
21         //
22         while (buffer.hasRemaining()) {
23             System.out.print(buffer.get());
24         }
25
26         //      clear()          position          0          limit
27         capacity
28         buffer.clear();
29         System.out.println("    clear()          ");
30         printState(buffer);
31     }
32
33
34     //      buffer capacity limit position mark
35     private static void printState(CharBuffer buffer) {
36         System.out.print("capacity: " + buffer.capacity());
37         System.out.print(", limit: " + buffer.limit());
38
39
40

```



```
        System.out.print(", position: " + buffer.position());
        System.out.print(", mark           : " + buffer.mark());
        System.out.println("\n");
    }
}
```

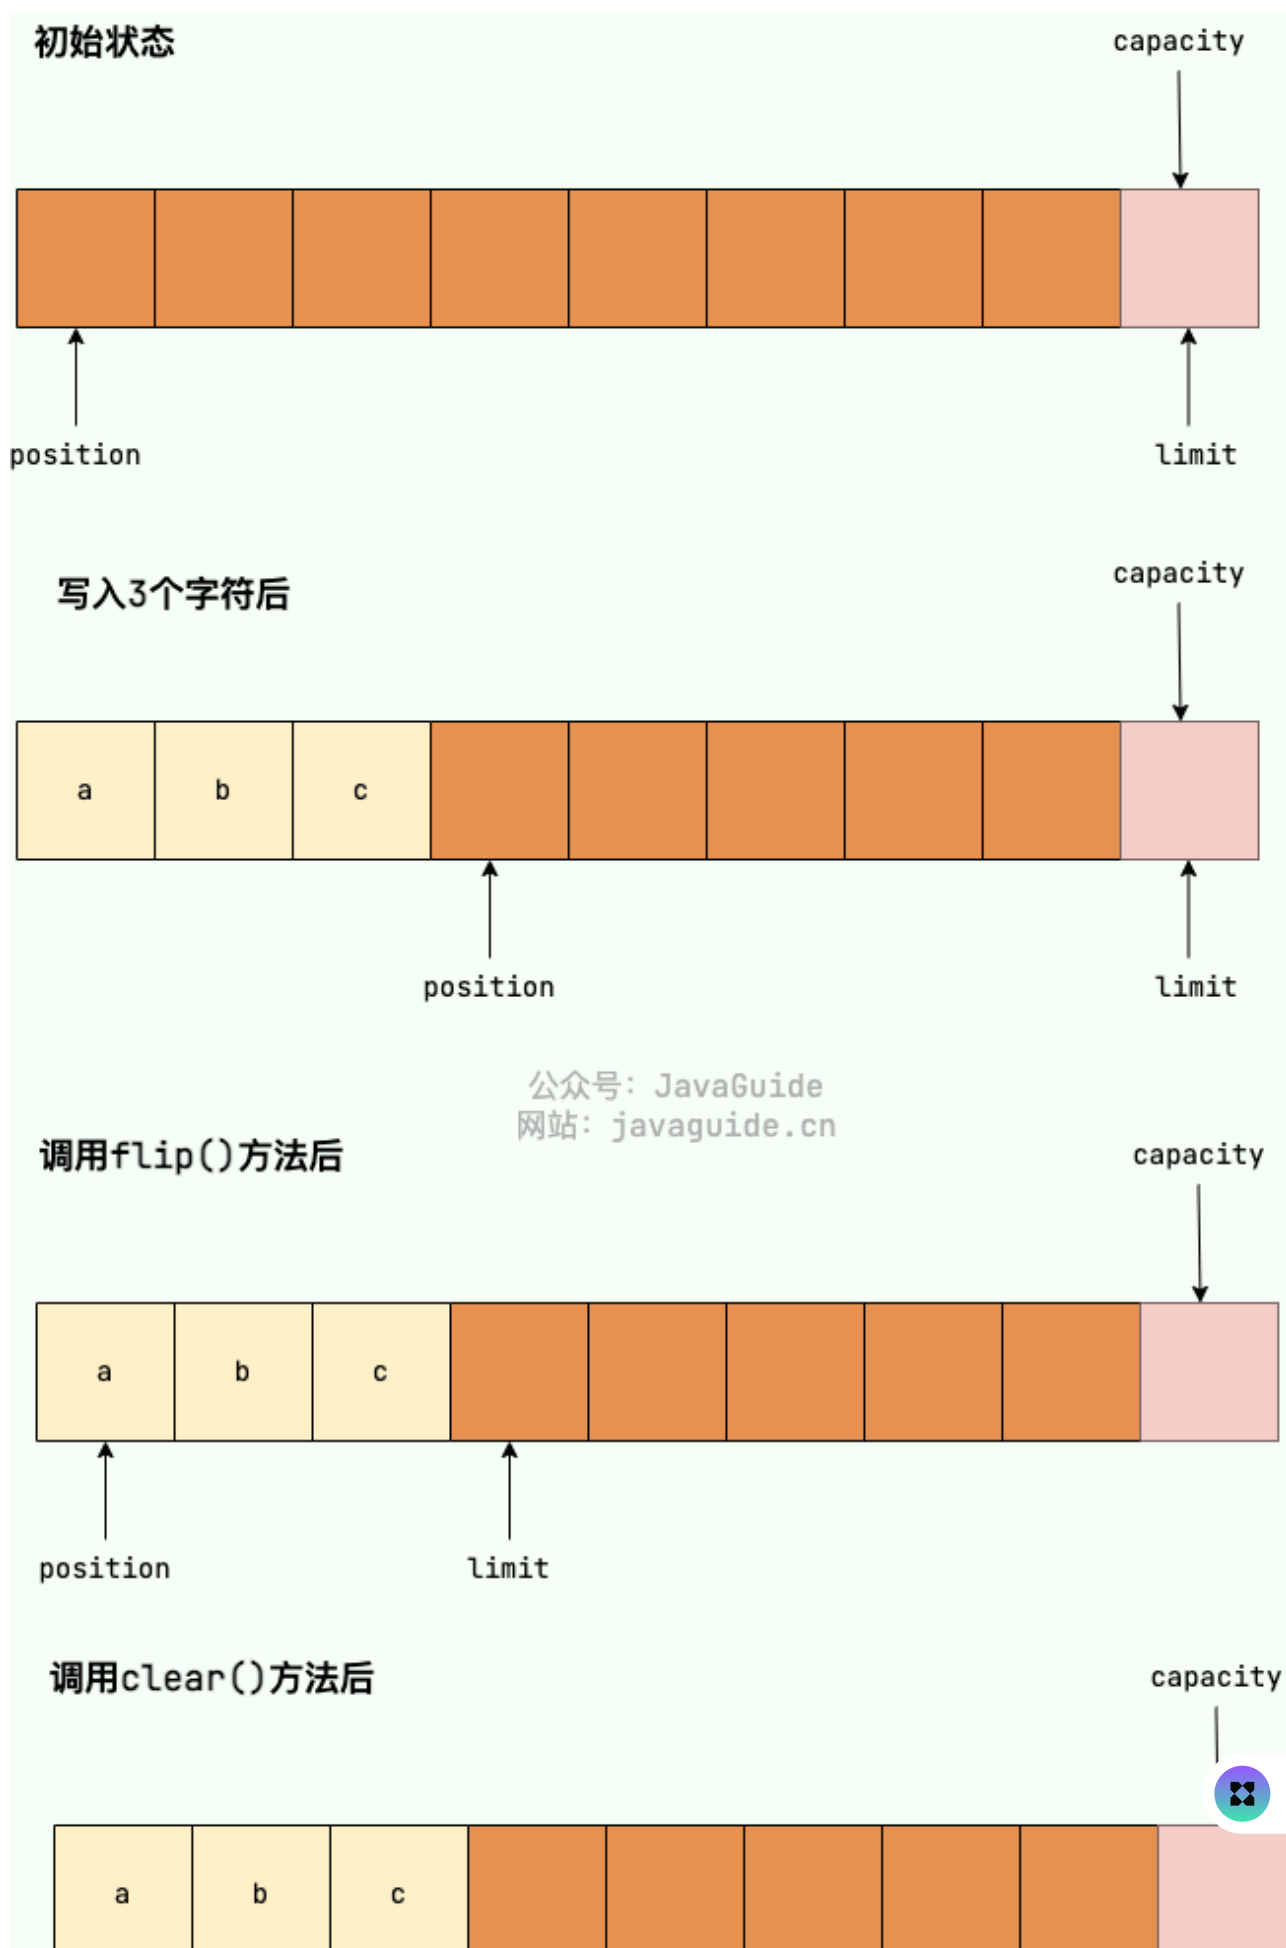
Output:

```
1
2 capacity: 8, limit: 8, position: 0
3
4     3
5 capacity: 8, limit: 8, position: 3
6
7     buffer
8
9     flip()
10 capacity: 8, limit: 3, position: 0
11
12     abc
13
14     clear()
15 capacity: 8, limit: 8, position: 0
```

bash

To help you understand, I drew a picture showing the process capacity and changes limit at position each stage.





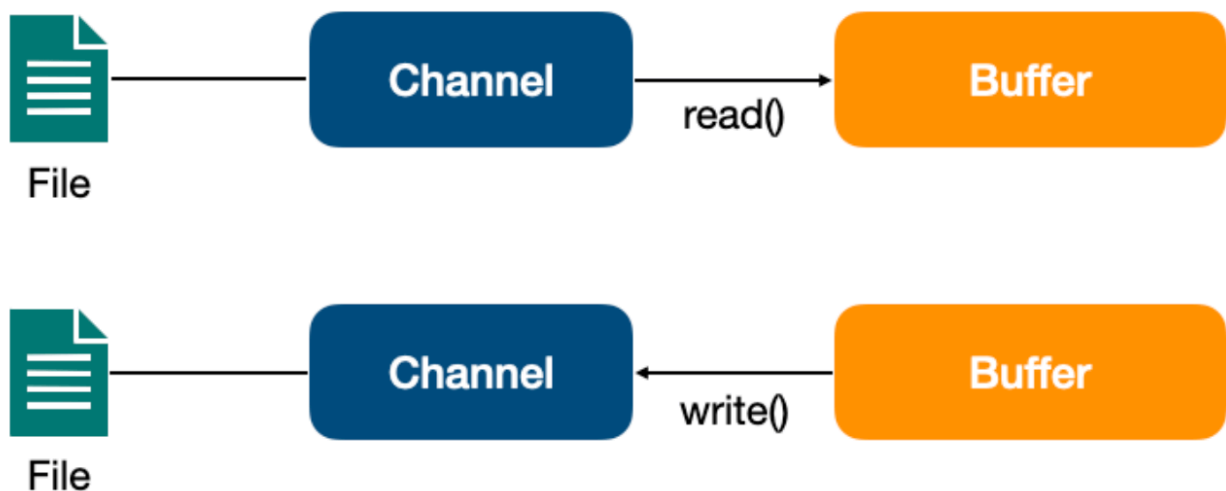


## Channel

A Channel is a channel that establishes a connection with a data source (such as a file, network socket, etc.). We can use it to read and write data, just like opening a water pipe and letting data flow freely in the Channel.

Streams in BIO are unidirectional and are classified as `InputStream` input streams and `OutputStream` output streams. Data is transmitted in one direction only. Channels differ from streams in that they are bidirectional and can be used for reading, writing, or both.

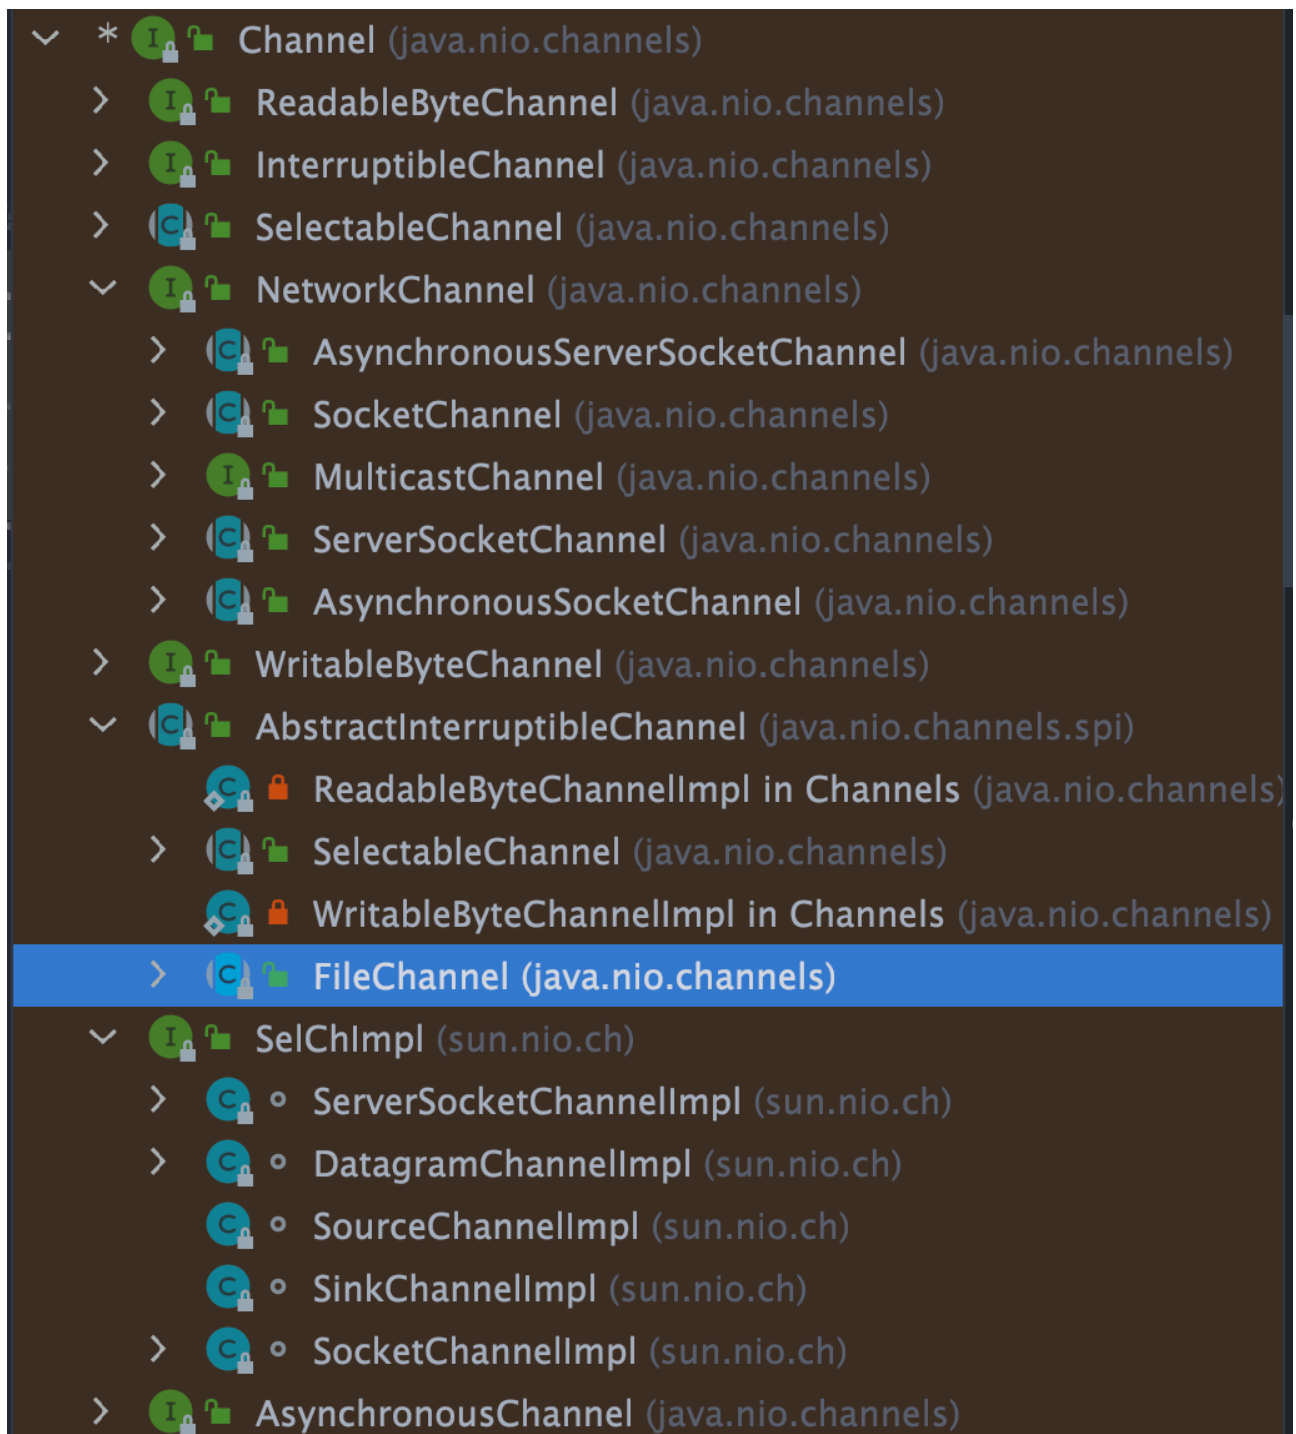
The Channel interacts with the Buffer introduced earlier. During a read operation, the data in the Channel is filled into the Buffer, while during a write operation, the data in the Buffer is written into the Channel.



In addition, because Channel is full-duplex, it can better map the API of the underlying operating system than streams. In particular, in the UNIX network programming model, the channels of the underlying operating system are full-duplex and support both read and write operations.

Channel The subclasses are shown below.

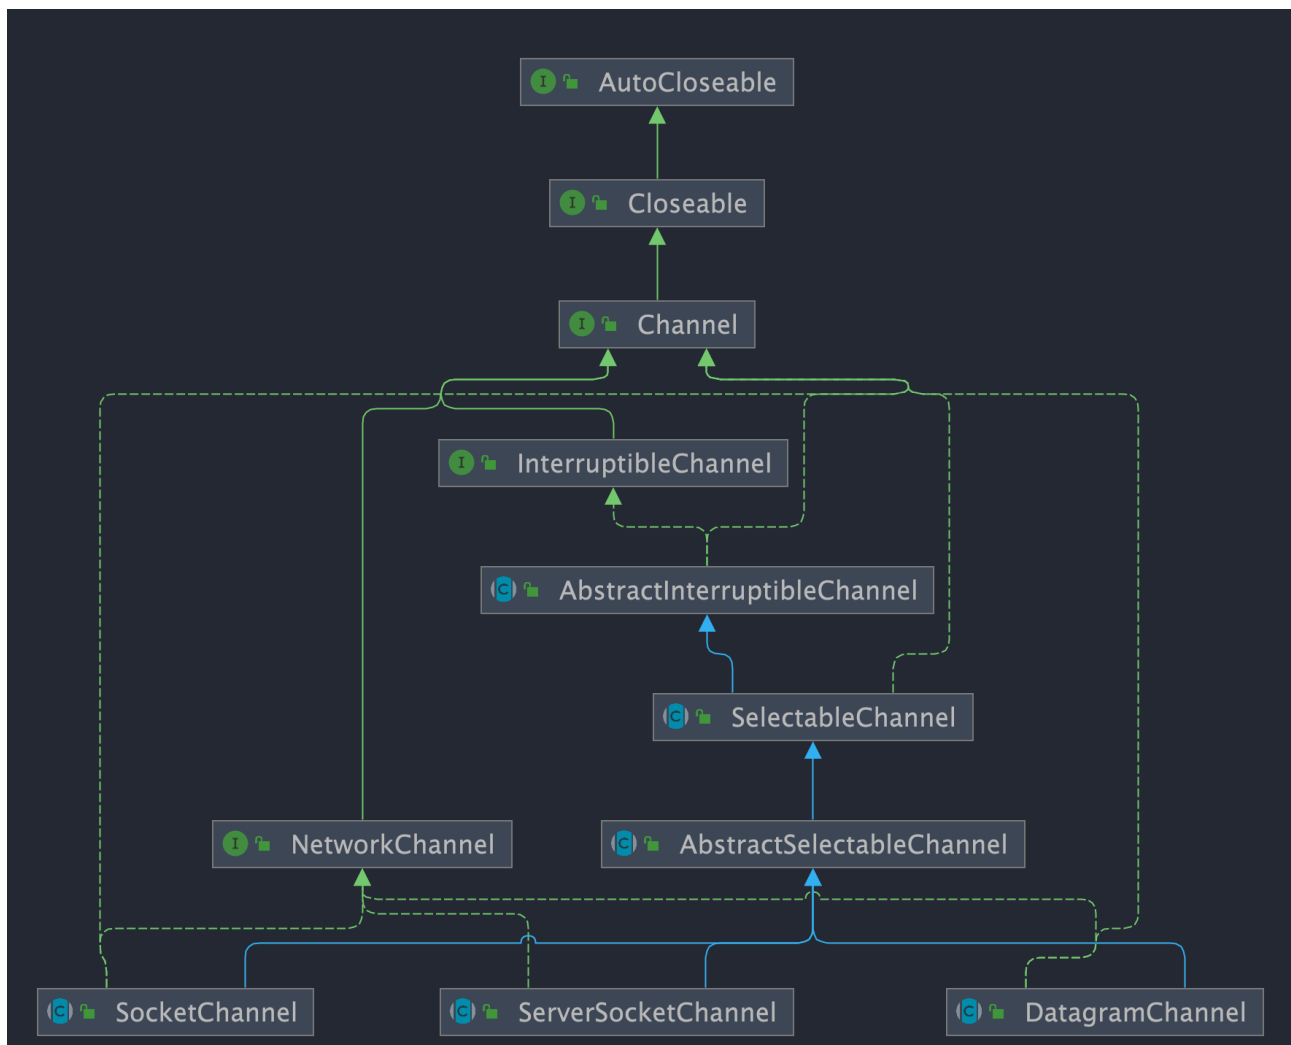




Among them, the most commonly used are the following types of channels:

- `FileChannel` : File access channel;
- `SocketChannel` , `ServerSocketChannel` : TCP communication channel;
- `DatagramChannel` : UDP communication channel;





The two core methods of Channel are:

1. read : Read data and write it into Buffer.
2. write : Write the data in the Buffer to the Channel.

Here we take as `FileChannel` an example to demonstrate how to read file data.

```

1 RandomAccessFile reader = new
2 RandomAccessFile("/Users/guide/Documents/test_read.in", "r"))
3 FileChannel channel = reader.getChannel();
4 ByteBuffer buffer = ByteBuffer.allocate(1024);
  channel.read(buffer);

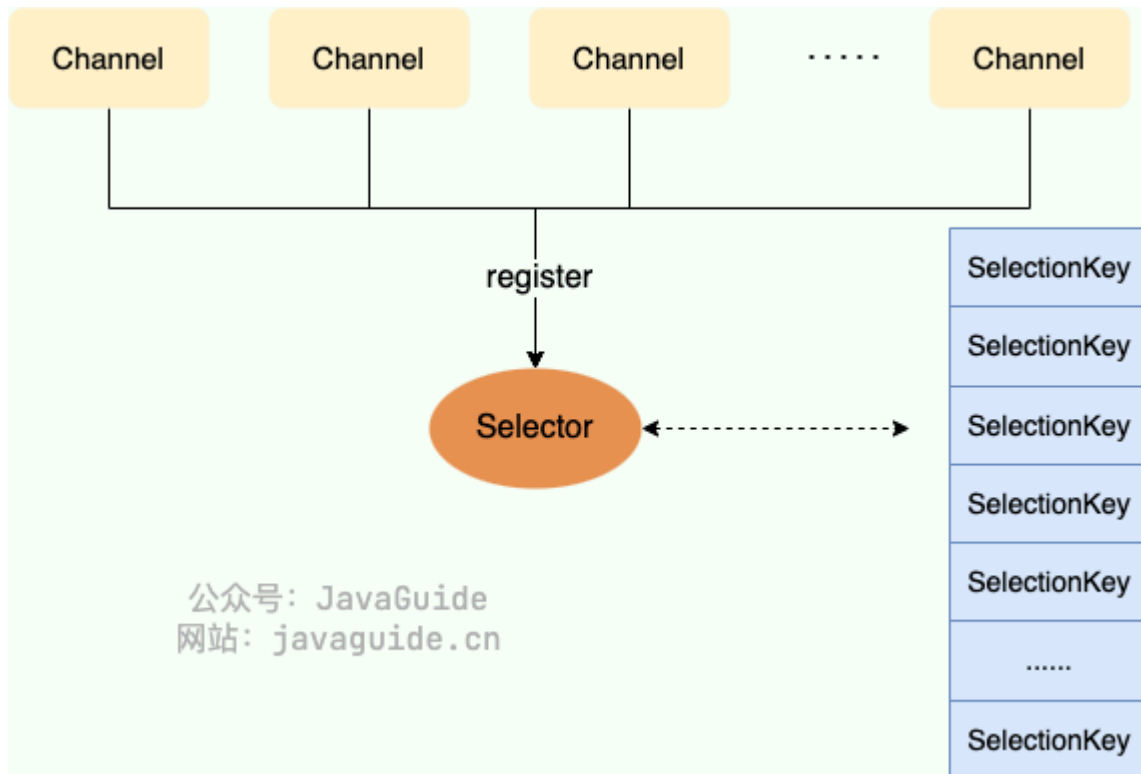
```

## Selector

The Selector is a key component in NIO, allowing a single thread to handle multiple Channels. The Selector is based on an event-driven I/O multiplexing model. Its main operating principle is: Channel events are registered through the Selector, and the Selector



continuously polls the registered Channels. When an event occurs, such as a new TCP connection, read, or write event on a Channel, the Channel enters the Ready state and is polled by the Selector. The Selector adds the relevant Channels to the Ready Set. The `SelectionKey` allows you to obtain a collection of Ready Channels, and then perform corresponding I/O operations on these Ready Channels.



A multiplexer Selector can poll multiple Channels simultaneously. Since the JDK uses `epoll()` instead of the traditional `select` implementation, it does not have a maximum connection handle 1024/2048 limit. This means that only one thread is needed to poll the Selector and connect to thousands of clients.

Selector can listen for the following four event types:

1. `SelectionKey.OP_ACCEPT` : Indicates that the channel accepts the connection event, which is usually used `ServerSocketChannel` .
2. `SelectionKey.OP_CONNECT` : Indicates that the channel has completed the connection event, which is usually used `SocketChannel` .
3. `SelectionKey.OP_READ` : Indicates that the channel is ready for reading, that is, there is data to be read.
4. `SelectionKey.OP_WRITE` : Indicates that the channel is ready for writing, that is, data can be written.



**Selector** It is an abstract class. You can create a Selector instance by calling its `open()` static methods. Selector can monitor the status `SelectableChannel` of multiple at the same time `IO` and is `IO` the core of non-blocking .

A Selector instance has three `SelectionKey` collections:

1. All `SelectionKey` collections: represents the ones registered on this Selector `Channel` . This collection can be `keys()` returned by the method.
2. The selected `SelectionKey` collection: represents all Channels that can `select()` be obtained through the method and need to be processed . This collection can be returned through `IO` `selectedKeys()`
3. Cancelled `SelectionKey` collection: represents all the canceled registration relationships `Channel` . When the method is executed next time `select()` , these `Channel` corresponding `SelectionKey` will be completely deleted. The program usually does not need to directly access this collection, and there is no exposed access method.

Here is a brief demonstration of how to traverse the selected `SelectionKey` collection and process it:

```

1  Set<SelectionKey> selectedKeys = selector.selectedKeys();
2  Iterator<SelectionKey> keyIterator = selectedKeys.iterator();
3  while (keyIterator.hasNext()) {
4      SelectionKey key = keyIterator.next();
5      if (key != null) {
6          if (key.isAcceptable()) {
7              // ServerSocketChannel
8          } else if (key.isConnectable()) {
9              //
10             } else if (key.isReadable()) {
11                 // Channel
12             } else if (key.isWritable()) {
13                 // Channel      Buffer
14             }
15         }
16         keyIterator.remove();
17     }

```

Selector also provides a series of `select()` related methods:

- `int select()` : Monitors all registered ones `Channel` . When there is an operation that needs to be processed among them `IO` , this method returns and



`SelectionKey` adds the corresponding to the selected `SelectionKey` set. This method returns `Channel` the number of these.

- `int select(long timeout)` : You can set the timeout period for `select()` the operation.
- `int selectNow()` : Executes an operation that returns immediately `select()` . Compared with the method without parameters `select()` , this method will not block the thread.
- `Selector wakeup()` : Make a method that has not yet returned `select()` return immediately.
- ...

A simple example of using `Selector` to implement network reading and writing:

```

1  import java.io.IOException;
2  import java.net.InetSocketAddress;
3  import java.nio.ByteBuffer;
4  import java.nio.channels.SelectionKey;
5  import java.nio.channels.Selector;
6  import java.nio.channels.ServerSocketChannel;
7  import java.nio.channels.SocketChannel;
8  import java.util.Iterator;
9  import java.util.Set;
10
11 public class NioSelectorExample {
12
13     public static void main(String[] args) {
14         try {
15             ServerSocketChannel serverSocketChannel =
16 ServerSocketChannel.open();
17             serverSocketChannel.configureBlocking(false);
18             serverSocketChannel.socket().bind(new
19 InetSocketAddress(8080));
20
21             Selector selector = Selector.open();
22             // ServerSocketChannel Selector OP_ACCEPT
23             serverSocketChannel.register(selector,
24 SelectionKey.OP_ACCEPT);
25
26             while (true) {
27                 int readyChannels = selector.select();
28
29                 if (readyChannels == 0) {

```

java



```

30         continue;
31     }
32
33     Set<SelectionKey> selectedKeys = selector.selectedKeys();
34     Iterator<SelectionKey> keyIterator =
35     selectedKeys.iterator();
36
37     while (keyIterator.hasNext()) {
38         SelectionKey key = keyIterator.next();
39
40         if (key.isAcceptable()) {
41             //
42             ServerSocketChannel server = (ServerSocketChannel)
43     key.channel();
44             SocketChannel client = server.accept();
45             client.configureBlocking(false);
46
47             // Selector OP_READ
48             client.register(selector, SelectionKey.OP_READ);
49         } else if (key.isReadable()) {
50             //
51             SocketChannel client = (SocketChannel) key.channel();
52             ByteBuffer buffer = ByteBuffer.allocate(1024);
53             int bytesRead = client.read(buffer);
54
55             if (bytesRead > 0) {
56                 buffer.flip();
57                 System.out.println(" " + new
58     String(buffer.array(), 0, bytesRead));
59                 // Selector OP_WRITE
60                 client.register(selector, SelectionKey.OP_WRITE);
61             } else if (bytesRead < 0) {
62                 //
63                 client.close();
64             }
65         } else if (key.isWritable()) {
66             //
67             SocketChannel client = (SocketChannel) key.channel();
68             ByteBuffer buffer = ByteBuffer.wrap("Hello,
69     Client!".getBytes());
70             client.write(buffer);
71
72             // Selector OP_READ

```



```
73         client.register(selector, SelectionKey.OP_READ);
74     }
75
76     keyIterator.remove();
    }
} catch (IOException e) {
    e.printStackTrace();
}
}
```

In this example, we create a simple server that listens on port 8080 and uses Selector to handle connection, read, and write events. When receiving data from the client, the server reads the data and prints it to the console, then replies "Hello, Client!" to the client.

## NIO Zero Copy

Zero copy is a common method to improve IO operation performance. Top open source projects such as ActiveMQ, Kafka, RocketMQ, QMQ, Netty, etc. all use zero copy.

Zero copy means that when a computer performs I/O operations, the CPU does not need to copy data from one storage area to another, thereby reducing context switches and CPU copy time. In other words, zero copy primarily addresses the problem of the operating system frequently copying data when processing I/O operations. Common zero copy implementation technologies include: `mmap+write`, `sendfile` and `sendfile + DMA gather copy`.

The following figure shows a comparison of various zero-copy technologies:

	CPU copy	DMA copy	System calls	Context Switching
Traditional methods	2	2	read+write	4
mmap+write	1	2	mmap+write	4
sendfile	1	2	sendfile	2



	<b>CPU copy</b>	<b>DMA copy</b>	<b>System calls</b>	<b>Context Switching</b>
sendfile + DMA gather copy	0	2	sendfile	2

As can be seen, both traditional I/O methods and those using zero-copy require two DMA (Direct Memory Access) copies. This is because both DMA operations rely on hardware. Zero-copy primarily reduces CPU copies and context switches.

Java supports zero copy:

- `MappedByteBuffer` NIO `mmap` provides a zero-copy implementation based on memory mapping (`map`). Under the hood, it actually invokes the Linux kernel's `mmap` system call. It can map a file or part of a file into memory, forming a virtual memory file. This allows direct access to data in memory without requiring system calls to read and write files.
- `FileChannel` This `transferTo()/transferFrom()` is an NIO `sendfile` implementation based on the zero-copy send file method (`sendFile()`). It actually uses a Linux kernel `sendfile` system call. It can send file data directly from disk to the network without going through a user-space buffer. For more information on `FileChannel` its usage, please refer to the article: [Java NIO FileChannel Usage](#).

Code example:

```

1  private void loadFileIntoMemory(File xmlFile) throws IOException {
2      FileInputStream fis = new FileInputStream(xmlFile);
3      // FileChannel
4      FileChannel fc = fis.getChannel();
5      // FileChannel.map()
6      MappedByteBuffer mmb = fc.map(FileChannel.MapMode.READ_ONLY, 0,
7      fc.size());
8      xmlFileBuffer = new byte[(int)fc.size()];
9      mmb.get(xmlFileBuffer);
10     fis.close();
11 }

```

## Summarize

In this article, we mainly introduced the core knowledge points of NIO, including NIO's core components and zero copy.



If you need to use NIO to build a network application, we recommend against using native NIO directly, as programming is complex and functionality is limited. Instead, we recommend using a mature NIO-based network programming framework, such as Netty. Netty offers several optimizations and extensions based on NIO, including support for multiple protocols and SSL/TLS.

## refer to

---

- A brief analysis of Java NIO: <https://tech.meituan.com/2016/11/04/nio.html>
- Interviewer: Do you know Java NIO? <https://mp.weixin.qq.com/s/mZobf-U8OSYQfHfyBEB6KA>
- Java NIO: Buffer, Channel and Selector: <https://www.javadoop.com/post/java-nio>

### JavaGuide官方公众号 (微信搜索JavaGuide)



- 1、公众号后台回复“PDF”获取原创PDF面试手册
- 2、公众号后台回复“学习路线”获取Java学习路线最新版
- 3、公众号后台回复“开源”获取优质Java开源项目合集
- 4、公众号后台回复“八股文”获取Java面试真题+面经

Recently Updated 2024/7/17 16:02

Contributors: dandelionstar , Guide , Mr.Hope , aucub , FengWei2000 , Guevara

---

Copyright © 2025 Guide

