

Java Stream Internals: Execution

Douglas C. Schmidt

d.schmidt@vanderbilt.edu

www.dre.vanderbilt.edu/~schmidt



Professor of Computer Science

**Institute for Software
Integrated Systems**

**Vanderbilt University
Nashville, Tennessee, USA**



Learning Objectives in this Part of the Lesson

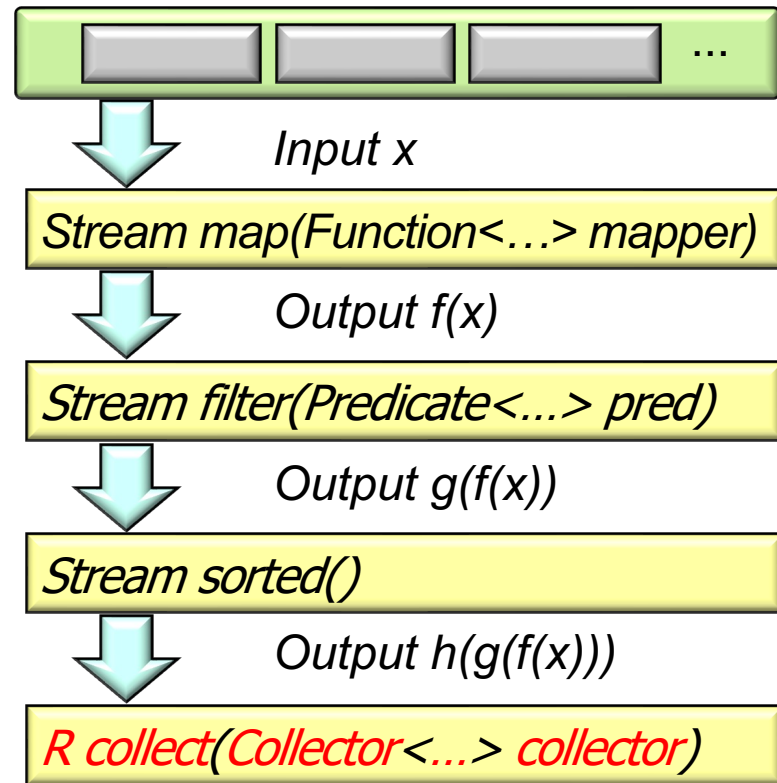
- Understand stream internals, e.g.
 - Know what can change & what can't
 - Recognize how a Java stream is constructed
- Be aware of how a Java stream is executed
 - e.g., how stateless & stateful intermediate operations & run-to-completion & short-circuiting terminal operations are run



Java Stream Execution

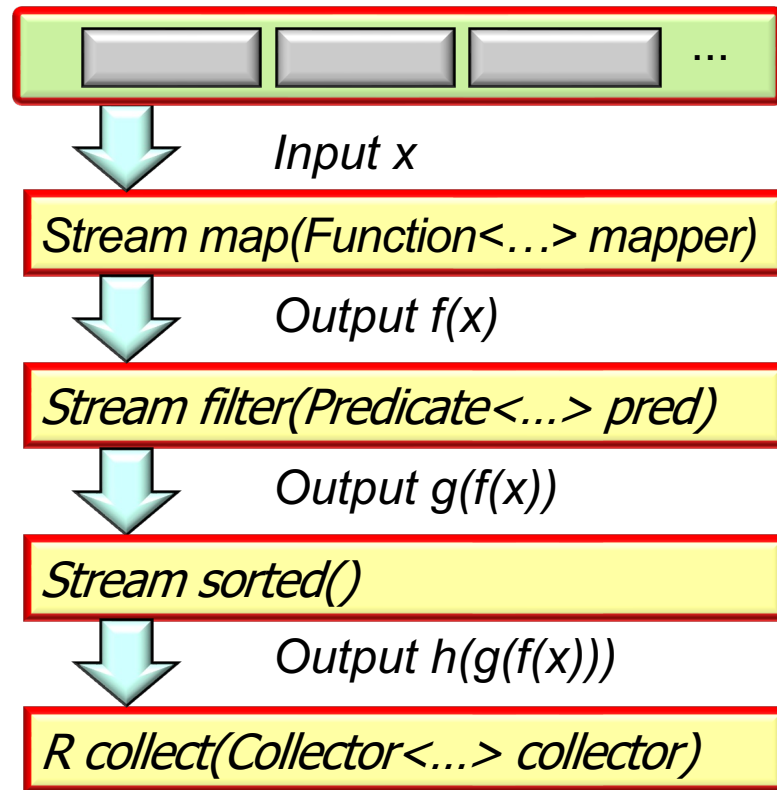
Java Stream Execution

- When terminal operation runs the streams framework picks an execution plan



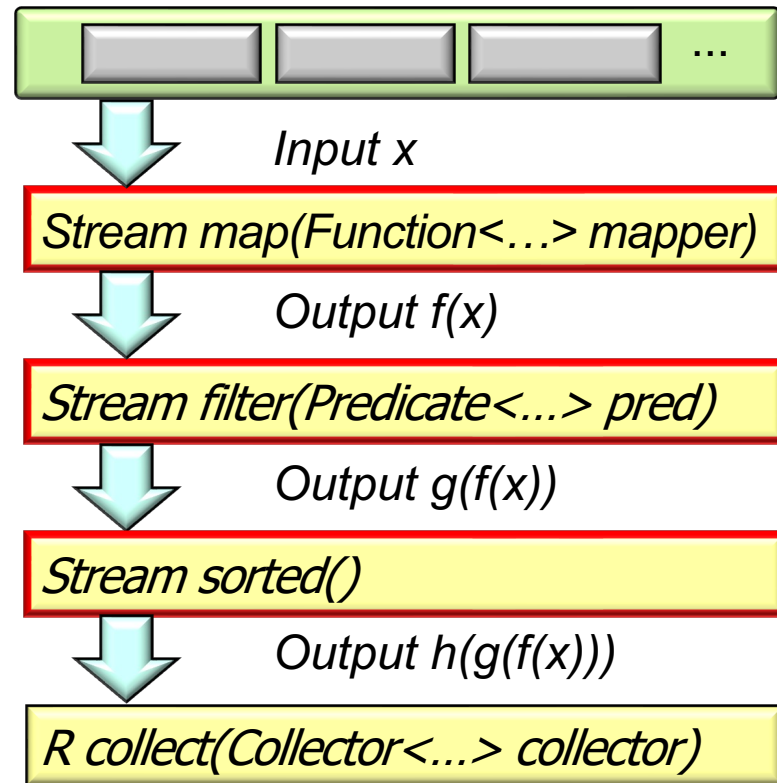
Java Stream Execution

- When terminal operation runs the streams framework picks an execution plan
- The plan is based on properties of the source & aggregate operations



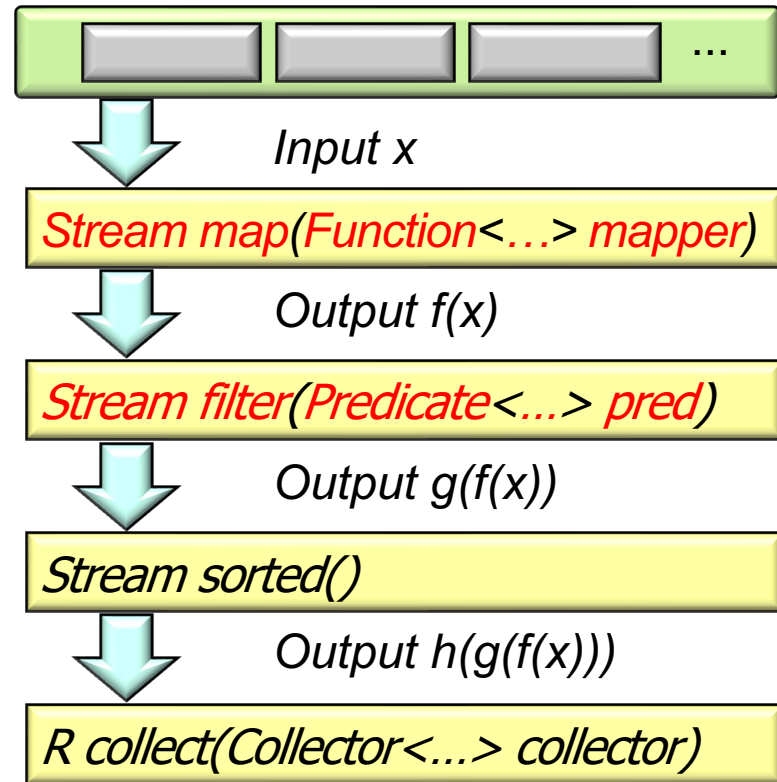
Java Stream Execution

- When terminal operation runs the streams framework picks an execution plan
 - The plan is based on properties of the source & aggregate operations
 - Intermediate operations are divided into two categories



Java Stream Execution

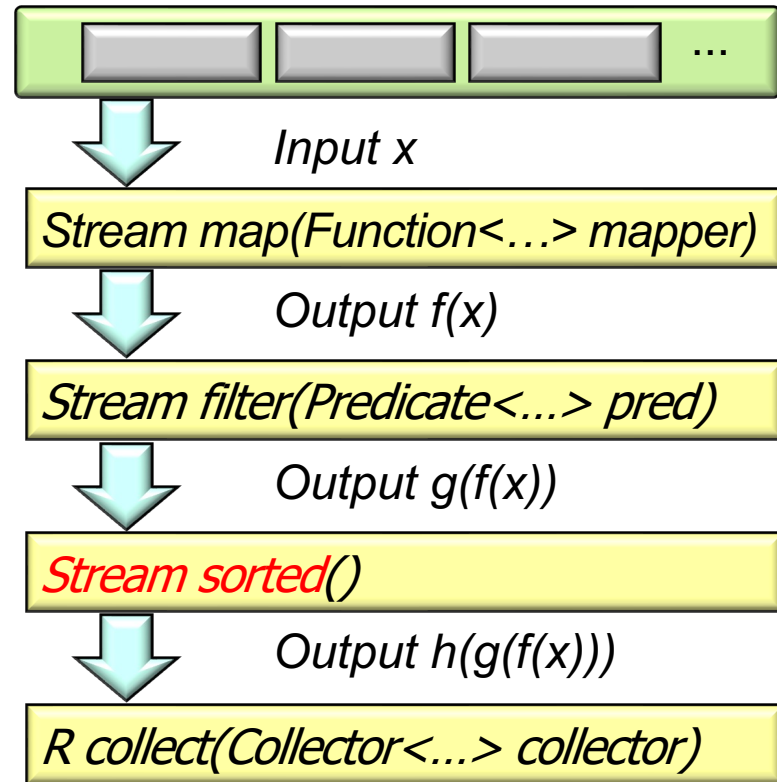
- When terminal operation runs the streams framework picks an execution plan
 - The plan is based on properties of the source & aggregate operations
 - Intermediate operations are divided into two categories:
 - Stateless
 - e.g., `filter()`, `map()`, `flatMap()`, `mapMulti()`, etc.



A pipeline with only stateless operations runs in one pass (even if it's parallel)

Java Stream Execution



- When terminal operation runs the streams framework picks an execution plan
 - The plan is based on properties of the source & aggregate operations
 - Intermediate operations are divided into two categories:
 - Stateless
 - Stateful
 - e.g., `sorted()`, `limit()`, `distinct()`, `dropWhile()`, `takeWhile()`, etc.



A pipeline with stateful operations is divided into sections & runs in multiple passes

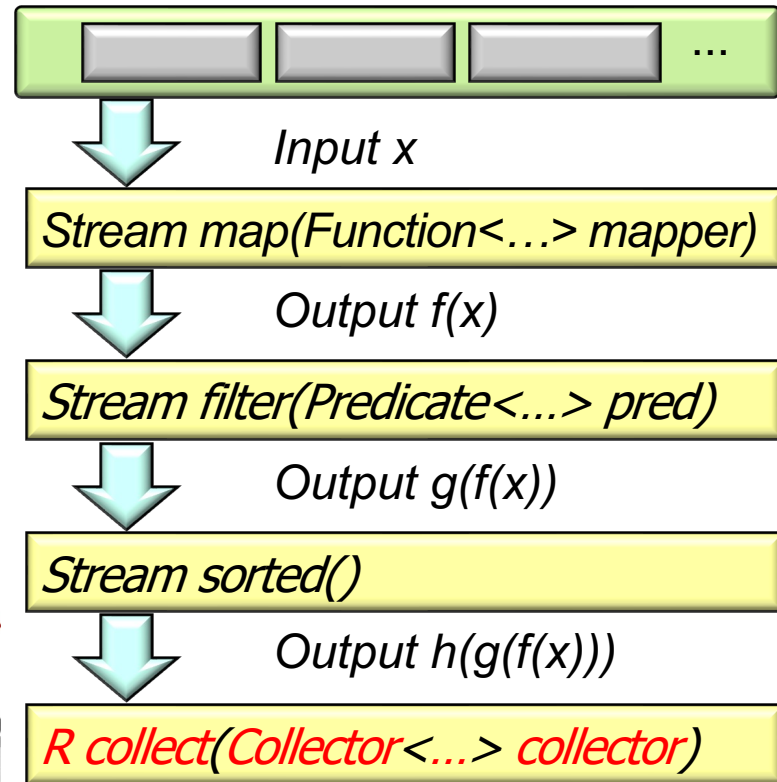
Java Stream Execution

- When terminal operation runs the streams framework picks an execution plan
 - The plan is based on properties of the source & aggregate operations
 - Intermediate operations are divided into two categories
 - Terminal operations are also divided into two categories



Java Stream Execution

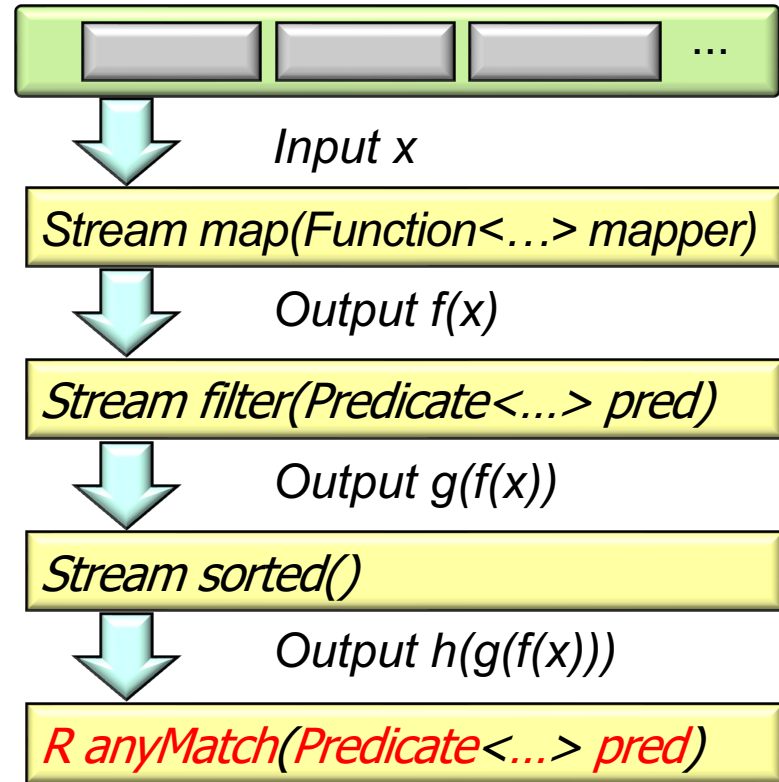
- When terminal operation runs the streams framework picks an execution plan
 - The plan is based on properties of the source & aggregate operations
 - Intermediate operations are divided into two categories
 - Terminal operations are also divided into two categories
 - Run-to-completion
 - e.g., `reduce()`, `collect()`, `forEach()`, etc.



These terminal operation process data in bulk using `Splitterator.forEachRemaining()`

Java Stream Execution

- When terminal operation runs the streams framework picks an execution plan
 - The plan is based on properties of the source & aggregate operations
 - Intermediate operations are divided into two categories
 - Terminal operations are also divided into two categories
 - Run-to-completion
 - Short-circuiting
 - e.g., `anyMatch()`, `findFirst()`, etc.



These terminal operation process data one element at a time using `tryAdvance()`.

End of Java Stream Internals: Execution