Memory Affinity in Multi-threading
The Bowtie2 Case Study

Claudia Misale, Marco Aldinucci
University of Torino, Italy

Massimo Torquati
University of Pisa, Italy

Memory bound problem

Synchronisation efficiency problem for very fine grain tasks (<1ms)

Load balancing problem for coarse grain tasks

Strongly lock-based due to shared data

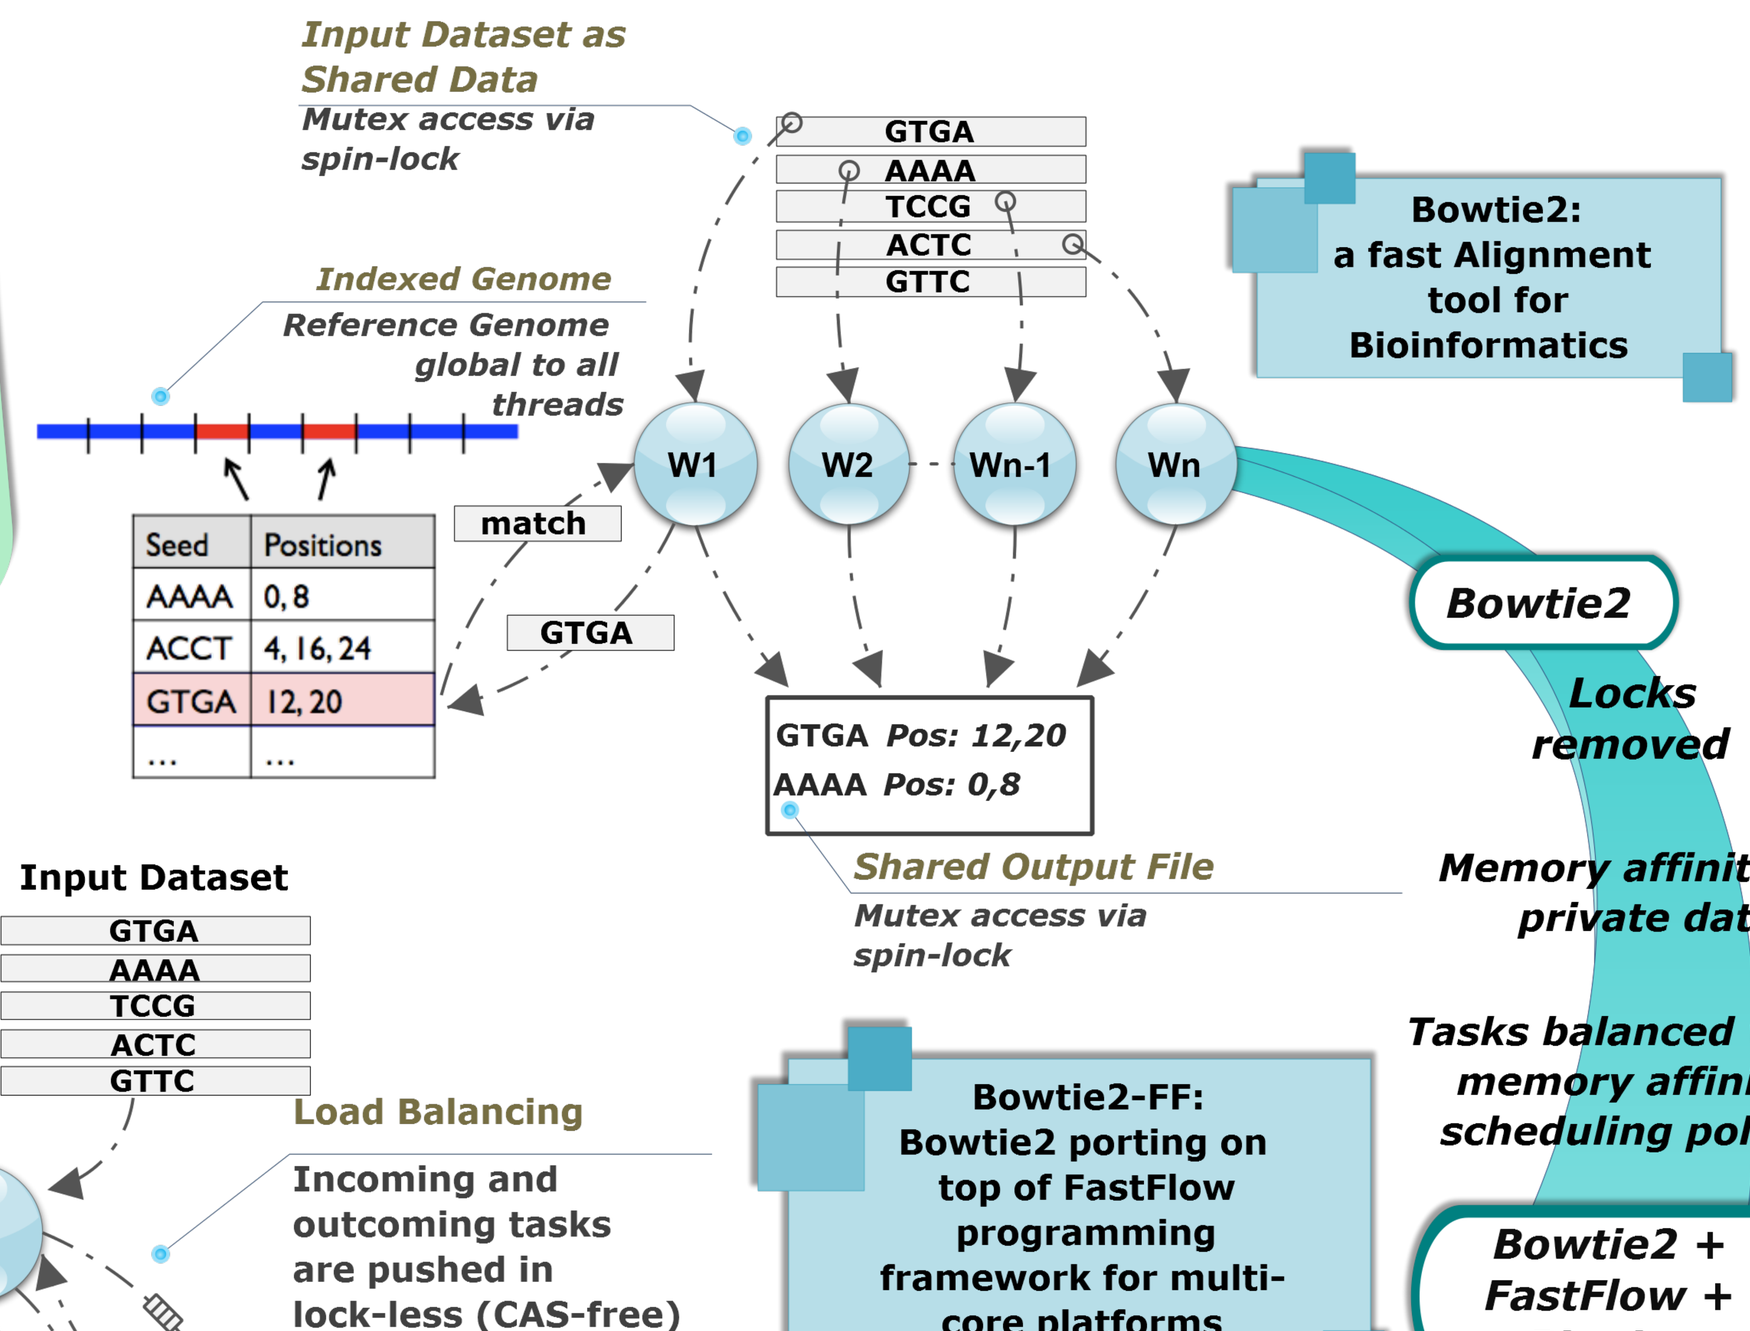

Bowtie2: a fast Alignment tool for Bioinformatics

Bowtie2-FF: Bowtie2 porting on top of FastFlow programming framework for multi-core platforms

Bowtie2

Locks removed

Memory affinity for private data

Tasks balanced with memory affinity scheduling policy

Bowtie2 + FastFlow + Pinning

Threads pinned on cores

Interleaving of shared data

Bowtie2 + FastFlow + Pinning + Interleaving

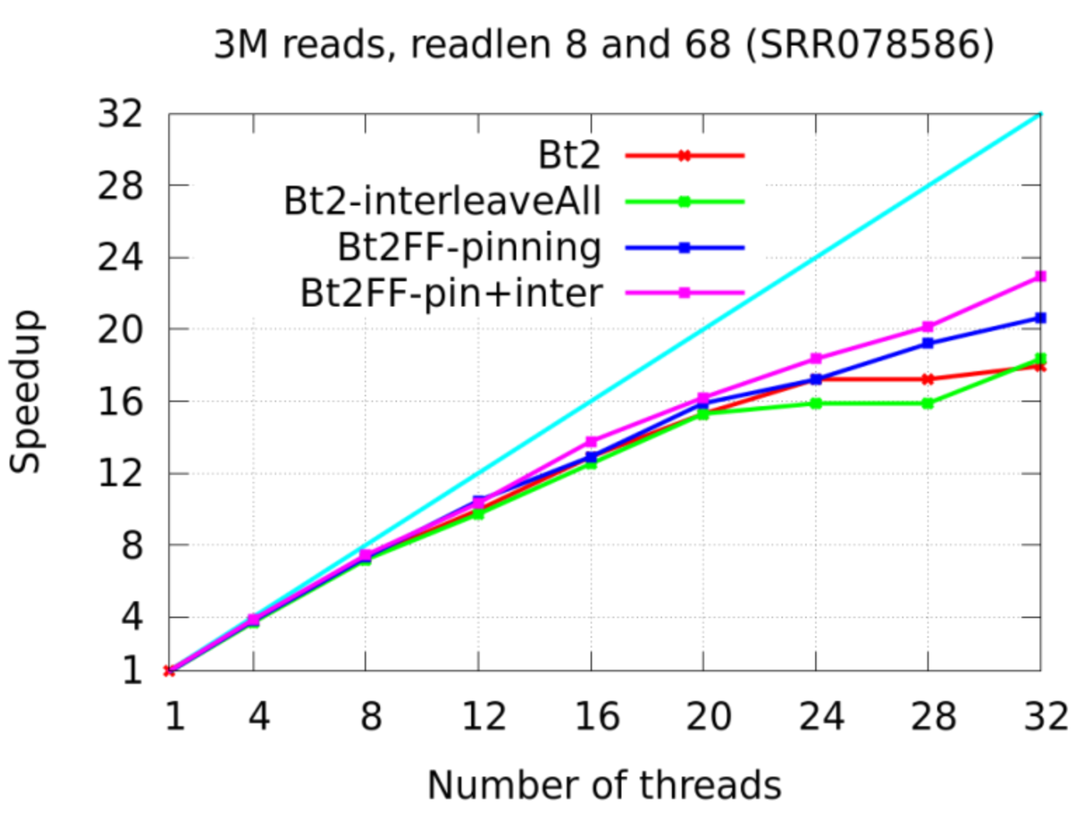

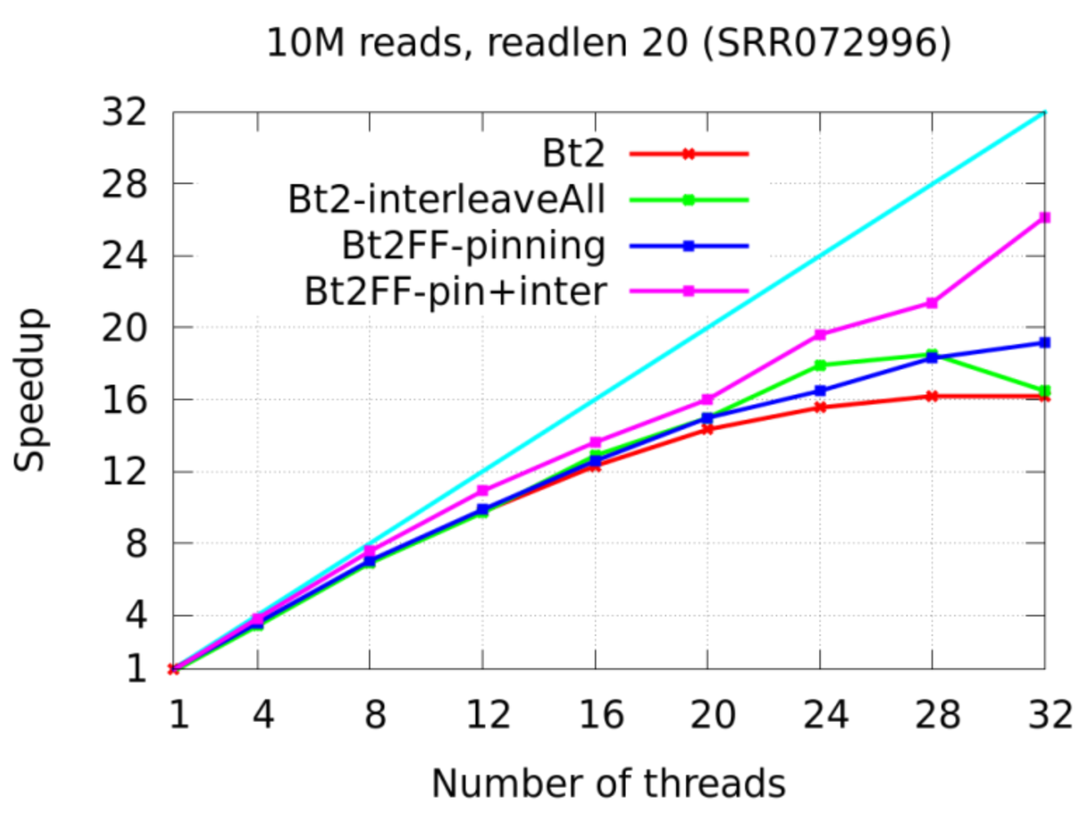

FastFlow web site: http://mc-fastflow.sourceforge.net/