

# Parallelism for the masses with modern C++

F. Giacomini

INFN-CNAF

Rimini, Workshop CCR, June 2018



# Proposal for Parallel Algorithms

## A Parallel Algorithms Library | N3554

Jared Hoberock    Jaydeep Marathe    Michael Garland    Olivier Giroux  
Vinod Grover    {jhoberock, jmarathe, mgarland, ogiroux, vgrover}@nvidia.com  
Artur Laksberg    Herb Sutter    {arturl, hsutter}@microsoft.com    Arch Robison  
arch.robison@intel.com

2013-03-15

“This proposal is motivated by a strong desire to provide a standard model of parallelism enabling **performance portability** across the broadest possible range of parallel architectures. For this reason, we have exposed parallelism in the most abstract manner possible in order to avoid presuming the existence of a particular parallel machine model.”

# C++ timeline




IS: trunk

TSes: feature branches for beta release & then merge



TS bars start and end where work on detailed specification wording starts ("adopt initial working draft") and ends ("send to publication")

Future starts/ends are shaded to indicate that dates, and TS branches are approximate and subject to change



“Talk is cheap, show me the code”

# The path to parallelism

Raw `for` loop → Standard algorithm → Parallel algorithm

# Execution policies

An execution policy indicates the kinds of parallelism allowed in the execution of an algorithm and expresses the consequent requirements on the element access functions

`seq` execution may not be parallelized

`par` execution may be parallelized

- access functions have to be thread-safe, i.e. no data race

`par_unseq` execution may be parallelized and vectorized

- access functions have to be thread- and vectorize-safe, e.g. no memory allocation and deallocation, no synchronization
- for GPUs