

CSE 633: Parallel Algorithms

Fall 2012

Parallelized Hash Collision Attacking

Benedikt Budig

Course Instructor:
Russ Miller

The Topic in a Nutshell



The Topic in a Nutshell



A *hash function* is a total function $h : \{0, 1\}^* \rightarrow \{0, 1\}^n$

The Topic in a Nutshell



A *hash function* is a total function $h : \{0, 1\}^* \rightarrow \{0, 1\}^n$ that maps arbitrarily long strings to strings of a fixed length n .

The Topic in a Nutshell



A *hash function* is a total function $h : \{0, 1\}^* \rightarrow \{0, 1\}^n$ that maps arbitrarily long strings to strings of a fixed length n .

Project Goal: Find Hash Collisions for given Hash

The Topic in a Nutshell



A *hash function* is a total function $h : \{0, 1\}^* \rightarrow \{0, 1\}^n$ that maps arbitrarily long strings to strings of a fixed length n .

Project Goal: Find Hash Collisions for given Hash

A *hash collision* occurs for two strings x, y if $h(x) = h(y)$,

The Topic in a Nutshell



A *hash function* is a total function $h : \{0, 1\}^* \rightarrow \{0, 1\}^n$ that maps arbitrarily long strings to strings of a fixed length n .

Project Goal: Find Hash Collisions for given Hash

A *hash collision* occurs for two strings x, y if $h(x) = h(y)$, that is, if h maps the two strings to the same hash value.

The Topic in a Nutshell



A *hash function* is a total function $h : \{0, 1\}^* \rightarrow \{0, 1\}^n$ that maps arbitrarily long strings to strings of a fixed length n .

Project Goal: Find Hash Collisions for given Hash

A *hash collision* occurs for two strings x, y if $h(x) = h(y)$, that is, if h maps the two strings to the same hash value.

—► find a string y such that $h(y)$ collides with given $h(x)$

The Topic in a Nutshell



A *hash function* is a total function $h : \{0, 1\}^* \rightarrow \{0, 1\}^n$ that maps arbitrarily long strings to strings of a fixed length n .

Project Goal: Find Hash Collisions for given Hash

A *hash collision* occurs for two strings x, y if $h(x) = h(y)$, that is, if h maps the two strings to the same hash value.

—► find a string y such that $h(y)$ collides with given $h(x)$

Reason to do that?

The Topic in a Nutshell



A *hash function* is a total function $h : \{0, 1\}^* \rightarrow \{0, 1\}^n$ that maps arbitrarily long strings to strings of a fixed length n .

Project Goal: Find Hash Collisions for given Hash

A *hash collision* occurs for two strings x, y if $h(x) = h(y)$, that is, if h maps the two strings to the same hash value.

—► find a string y such that $h(y)$ collides with given $h(x)$

Reason to do that: Cryptographic Applications



The Topic in a Nutshell



A *hash function* is a total function $h : \{0, 1\}^* \rightarrow \{0, 1\}^n$ that maps arbitrarily long strings to strings of a fixed length n .

Project Goal: Find Hash Collisions for given Hash

A *hash collision* occurs for two strings x, y if $h(x) = h(y)$, that is, if h maps the two strings to the same hash value.

—► find a string y such that $h(y)$ collides with given $h(x)$

Reason to do that: Cryptographic Applications

- secure storage of passwords



The Topic in a Nutshell



A *hash function* is a total function $h : \{0, 1\}^* \rightarrow \{0, 1\}^n$ that maps arbitrarily long strings to strings of a fixed length n .

Project Goal: Find Hash Collisions for given Hash

A *hash collision* occurs for two strings x, y if $h(x) = h(y)$, that is, if h maps the two strings to the same hash value.

—► find a string y such that $h(y)$ collides with given $h(x)$

Reason to do that: Cryptographic Applications

- secure storage of passwords
- digital signature schemes



The Topic in a Nutshell



A *hash function* is a total function $h : \{0, 1\}^* \rightarrow \{0, 1\}^n$ that maps arbitrarily long strings to strings of a fixed length n .

Project Goal: Find Hash Collisions for given Hash

A *hash collision* occurs for two strings x, y if $h(x) = h(y)$, that is, if h maps the two strings to the same hash value.

—► find a string y such that $h(y)$ collides with given $h(x)$

Reason to do that: Cryptographic Applications

- secure storage of passwords
- digital signature schemes

We focus on
—► MD5



Parallel Approach

Parallel Approach

Input:

Parallel Approach

Input: hash $\text{md5}(x)$ of unknown string x of length $|x| \leq n$

Parallel Approach

Input: hash $\text{md5}(x)$ of unknown string x of length $|x| \leq n$

Output:

Parallel Approach

Input: hash $\text{md5}(x)$ of unknown string x of length $|x| \leq n$

Output: string y such that $\text{md5}(y) = \text{md5}(x)$ and $|y| \leq n$

Parallel Approach

Input: hash $\text{md5}(x)$ of unknown string x of length $|x| \leq n$

Output: string y such that $\text{md5}(y) = \text{md5}(x)$ and $|y| \leq n$

Parallel approach

Parallel Approach

Input: hash $\text{md5}(x)$ of unknown string x of length $|x| \leq n$

Output: string y such that $\text{md5}(y) = \text{md5}(x)$ and $|y| \leq n$

Parallel approach

1. based on their ID, m parallel processes take a subset of the possible strings $\{0, 1\}^{\leq n}$

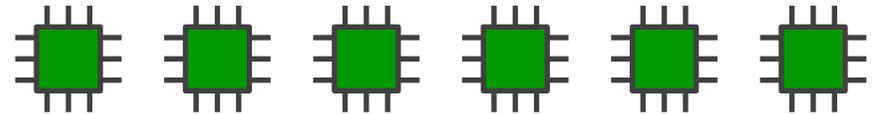
Parallel Approach

Input: hash $\text{md5}(x)$ of unknown string x of length $|x| \leq n$

Output: string y such that $\text{md5}(y) = \text{md5}(x)$ and $|y| \leq n$

Parallel approach

1. based on their ID, m parallel processes take a subset of the possible strings $\{0, 1\}^{\leq n}$



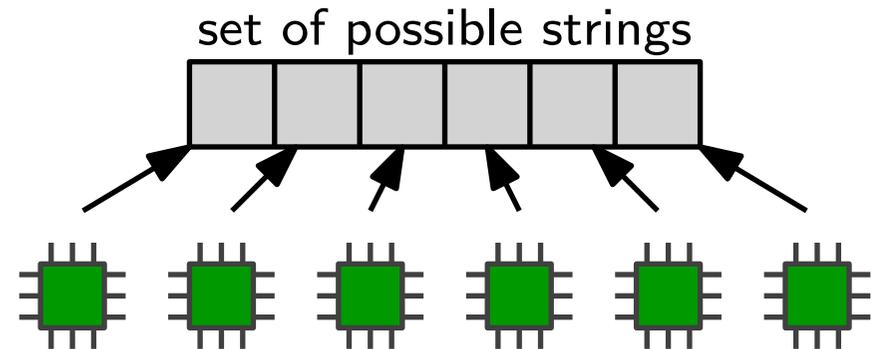
Parallel Approach

Input: hash $\text{md5}(x)$ of unknown string x of length $|x| \leq n$

Output: string y such that $\text{md5}(y) = \text{md5}(x)$ and $|y| \leq n$

Parallel approach

1. based on their ID, m parallel processes take a subset of the possible strings $\{0, 1\}^{\leq n}$



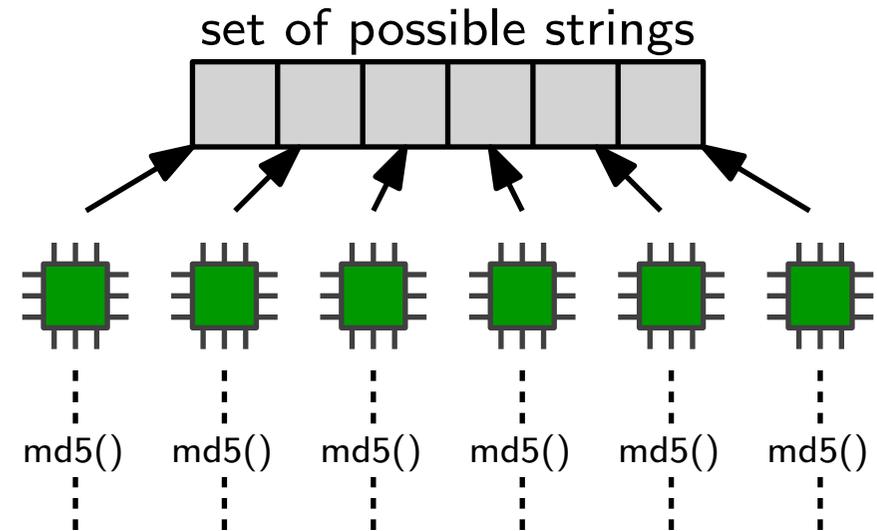
Parallel Approach

Input: hash $\text{md5}(x)$ of unknown string x of length $|x| \leq n$

Output: string y such that $\text{md5}(y) = \text{md5}(x)$ and $|y| \leq n$

Parallel approach

1. based on their ID, m parallel processes take a subset of the possible strings $\{0, 1\}^{\leq n}$
2. each process calculates hashes of the strings assigned to it



Parallel Approach

Input: hash $\text{md5}(x)$ of unknown string x of length $|x| \leq n$

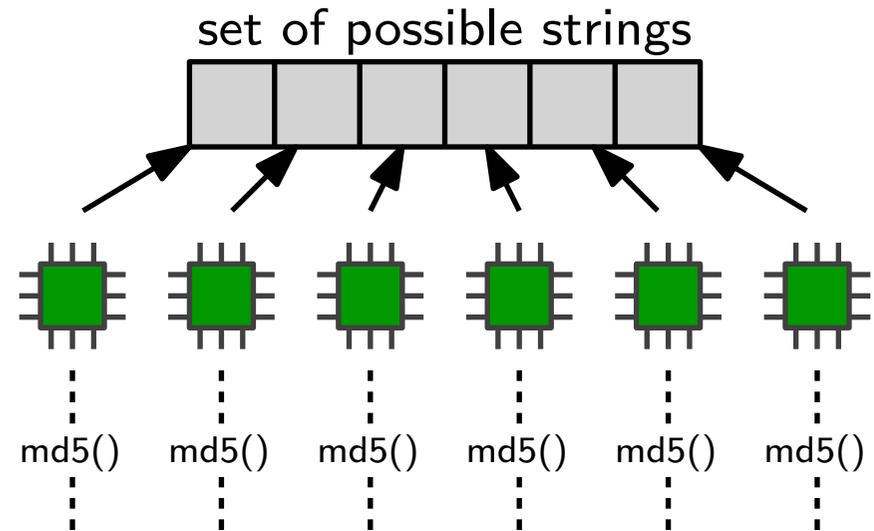
Output: string y such that $\text{md5}(y) = \text{md5}(x)$ and $|y| \leq n$

Parallel approach

1. based on their ID, m parallel processes take a subset of the possible strings $\{0, 1\}^{\leq n}$

2. each process calculates hashes of the strings assigned to it

3. as soon as one process calculates a hash equal to the input hash, all processes terminate



Parallel Approach

Input: hash $\text{md5}(x)$ of unknown string x of length $|x| \leq n$

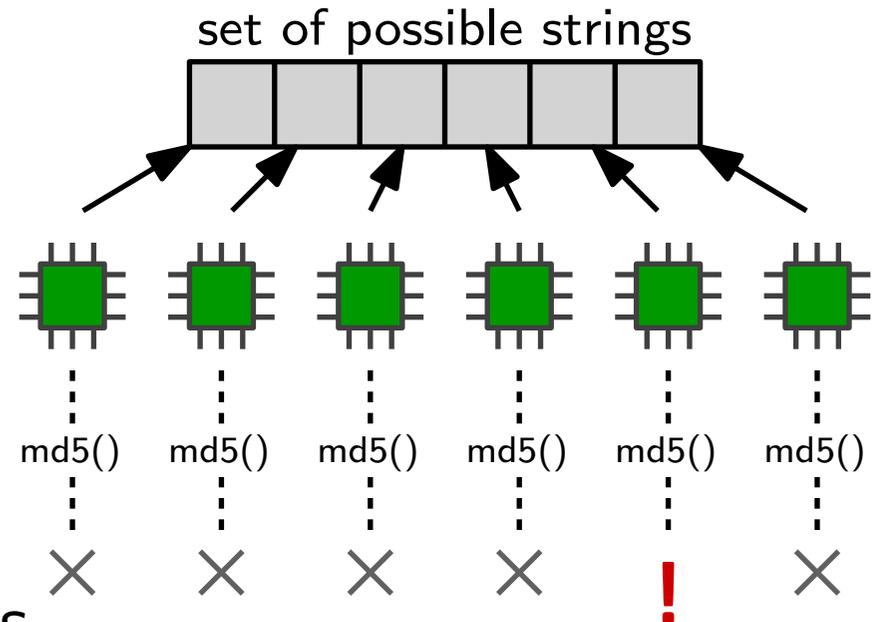
Output: string y such that $\text{md5}(y) = \text{md5}(x)$ and $|y| \leq n$

Parallel approach

1. based on their ID, m parallel processes take a subset of the possible strings $\{0, 1\}^{\leq n}$

2. each process calculates hashes of the strings assigned to it

3. as soon as one process calculates a hash equal to the input hash, all processes terminate



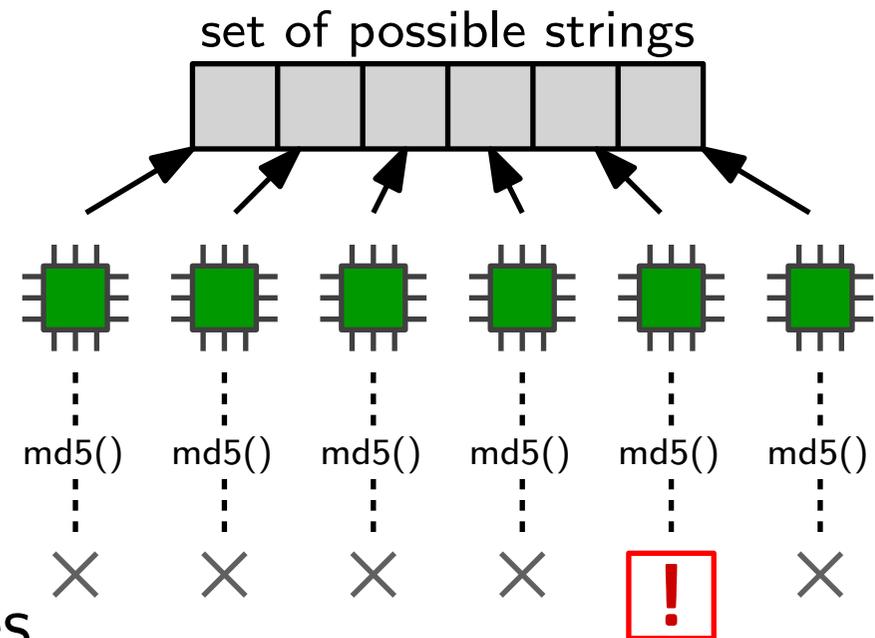
Parallel Approach

Input: hash $\text{md5}(x)$ of unknown string x of length $|x| \leq n$

Output: string y such that $\text{md5}(y) = \text{md5}(x)$ and $|y| \leq n$

Parallel approach

1. based on their ID, m parallel processes take a subset of the possible strings $\{0, 1\}^{\leq n}$
2. each process calculates hashes of the strings assigned to it
3. as soon as one process calculates a hash equal to the input hash, all processes terminate
4. return the colliding string



Technical Realization

Technical Realization

Hardware



image source: CCR website

Technical Realization

Hardware

- use of CPUs



image source: CCR website

Technical Realization

Hardware

- use of CPUs
- use of a multi-core system



image source: CCR website

Technical Realization

Hardware

- use of CPUs
- use of a multi-core system
- Infiniband network



image source: CCR website

Technical Realization

Hardware

- use of CPUs
- use of a multi-core system
- Infiniband network
- tests on CCR machines with 12 cores and 32 cores



image source: CCR website

Technical Realization

Hardware

- use of CPUs
- use of a multi-core system
- Infiniband network
- tests on CCR machines with 12 cores and 32 cores



image source: CCR website

Software and Implementation

Technical Realization

Hardware

- use of CPUs
- use of a multi-core system
- Infiniband network
- tests on CCR machines with 12 cores and 32 cores

Software and Implementation

- implementation using the C++ programming language



image source: CCR website

Technical Realization

Hardware

- use of CPUs
- use of a multi-core system
- Infiniband network
- tests on CCR machines with 12 cores and 32 cores



image source: CCR website

Software and Implementation

- implementation using the C++ programming language
- one implementation using OpenMP

Technical Realization

Hardware

- use of CPUs
- use of a multi-core system
- Infiniband network
- tests on CCR machines with 12 cores and 32 cores



image source: CCR website

Software and Implementation

- implementation using the C++ programming language
- one implementation using OpenMP
- and another version using MPI

Technical Realization

Hardware

- use of CPUs
- use of a multi-core system
- Infiniband network
- tests on CCR machines with 12 cores and 32 cores



image source: CCR website

Software and Implementation

- implementation using the C++ programming language
- one implementation using OpenMP
- and another version using MPI

Future Work

Technical Realization

Hardware

- use of CPUs
- use of a multi-core system
- Infiniband network
- tests on CCR machines with 12 cores and 32 cores



image source: CCR website

Software and Implementation

- implementation using the C++ programming language
- one implementation using OpenMP
- and another version using MPI

Future Work

- tweaks for MPI and OpenMPI

Technical Realization

Hardware

- use of CPUs
- use of a multi-core system
- Infiniband network
- tests on CCR machines with 12 cores and 32 cores



image source: CCR website

Software and Implementation

- implementation using the C++ programming language
- one implementation using OpenMP
- and another version using MPI

Future Work

- tweaks for MPI and OpenMPI
- implementation using CUDA

Benchmarks



Benchmarks

First Test: OpenMP on 12 Core System



Benchmarks

First Test: OpenMP on 12 Core System

- 12 Intel Xeon E5645 at 2.40GHz



Benchmarks

First Test: OpenMP on 12 Core System

- 12 Intel Xeon E5645 at 2.40GHz
- input: $\text{md5}(x)$ with $x \in \{0, 1\}^{24}$

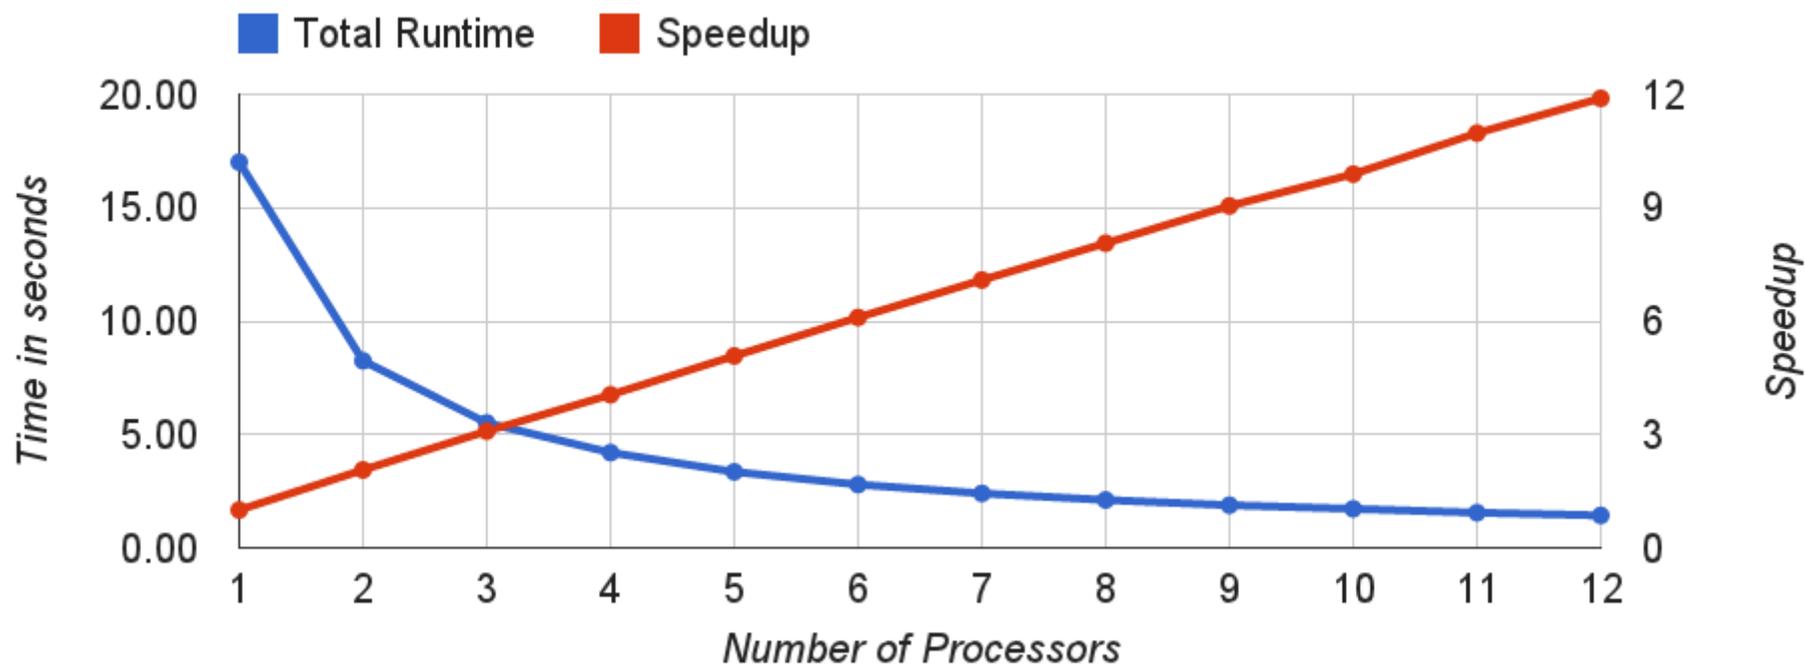


Benchmarks

First Test: OpenMP on 12 Core System



- 12 Intel Xeon E5645 at 2.40GHz
- input: $\text{md5}(x)$ with $x \in \{0, 1\}^{24}$

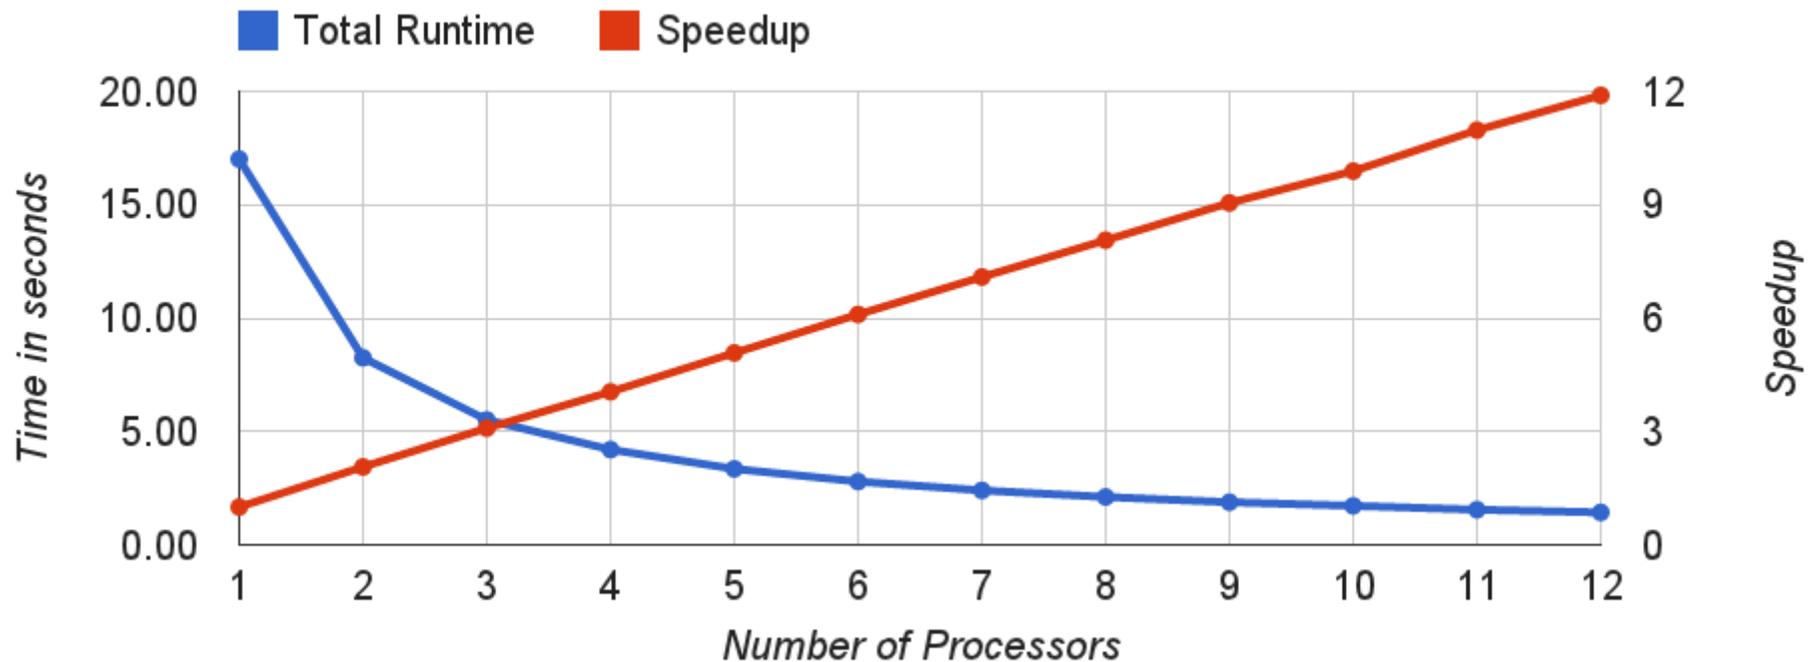


Benchmarks

First Test: OpenMP on 12 Core System



- 12 Intel Xeon E5645 at 2.40GHz
- input: md5(x) with $x \in \{0, 1\}^{24}$



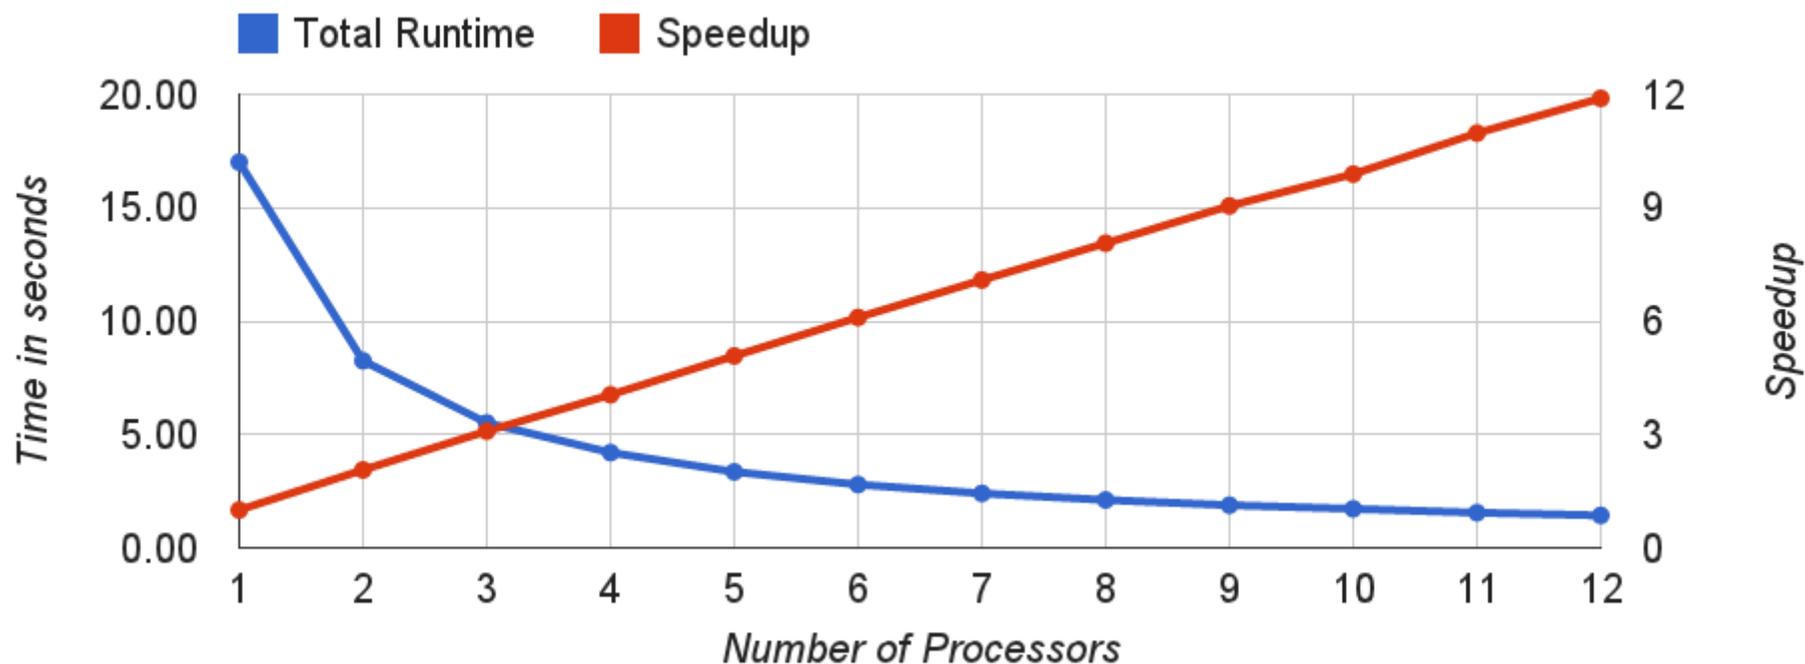
- 1.43 seconds to find collision

Benchmarks

First Test: OpenMP on 12 Core System



- 12 Intel Xeon E5645 at 2.40GHz
- input: $\text{md5}(x)$ with $x \in \{0, 1\}^{24}$



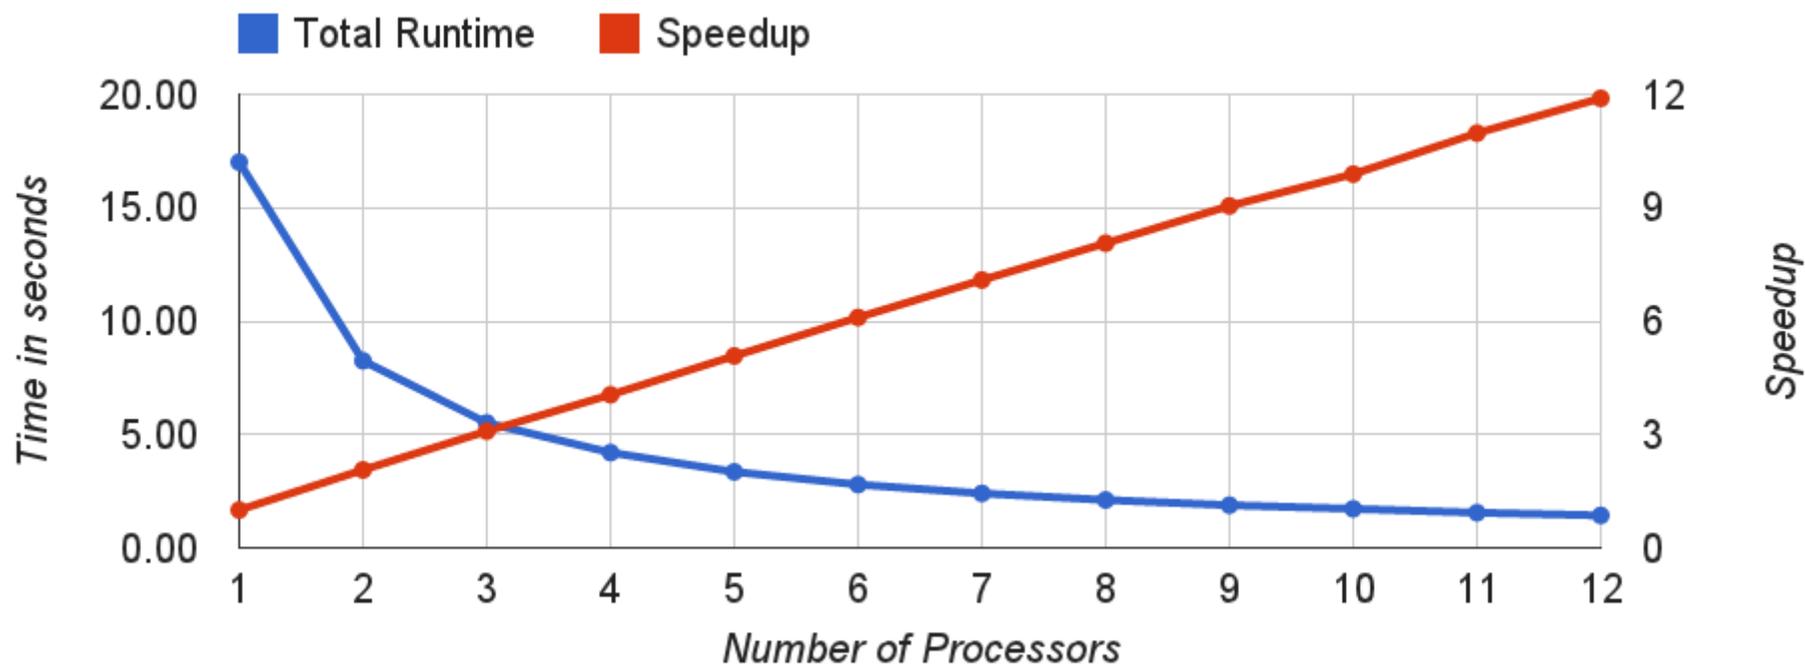
- 1.43 seconds to find collision
- approx. 4 million strings tried

Benchmarks



First Test: OpenMP on 12 Core System

- 12 Intel Xeon E5645 at 2.40GHz
- input: $\text{md5}(x)$ with $x \in \{0, 1\}^{24}$



- 1.43 seconds to find collision
- approx. 4 million strings tried

→ linear speedup

Benchmarks



Benchmarks

Second Test: OpenMP on 12 Core System



Benchmarks

Second Test: OpenMP on 12 Core System

- 12 Intel Xeon E5645 at 2.40GHz



Benchmarks

Second Test: OpenMP on 12 Core System



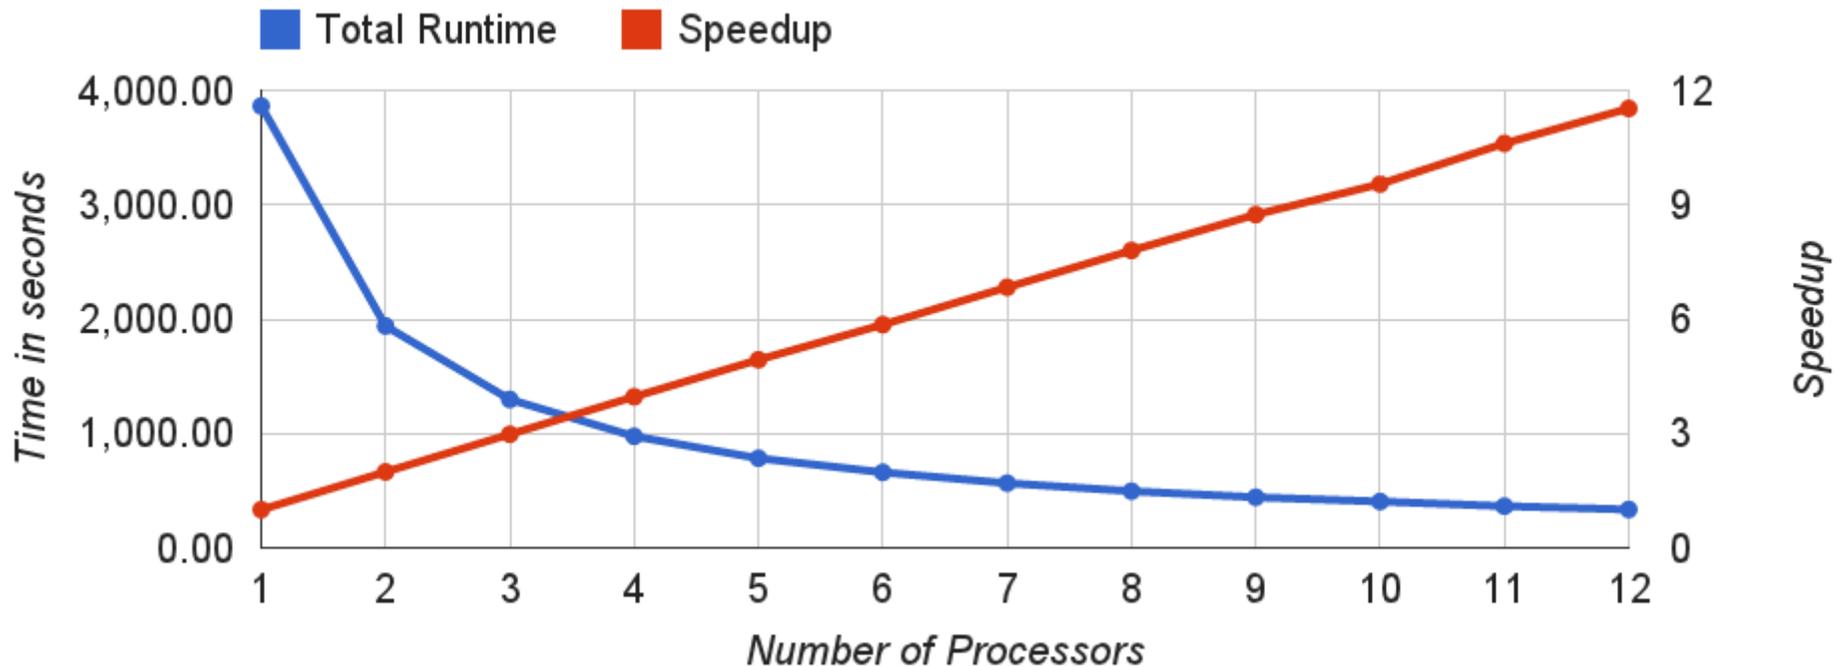
- 12 Intel Xeon E5645 at 2.40GHz
- input: $\text{md5}(x)$ with $x \in \{0, 1\}^{32}$

Benchmarks

Second Test: OpenMP on 12 Core System



- 12 Intel Xeon E5645 at 2.40GHz
- input: md5(x) with $x \in \{0, 1\}^{32}$

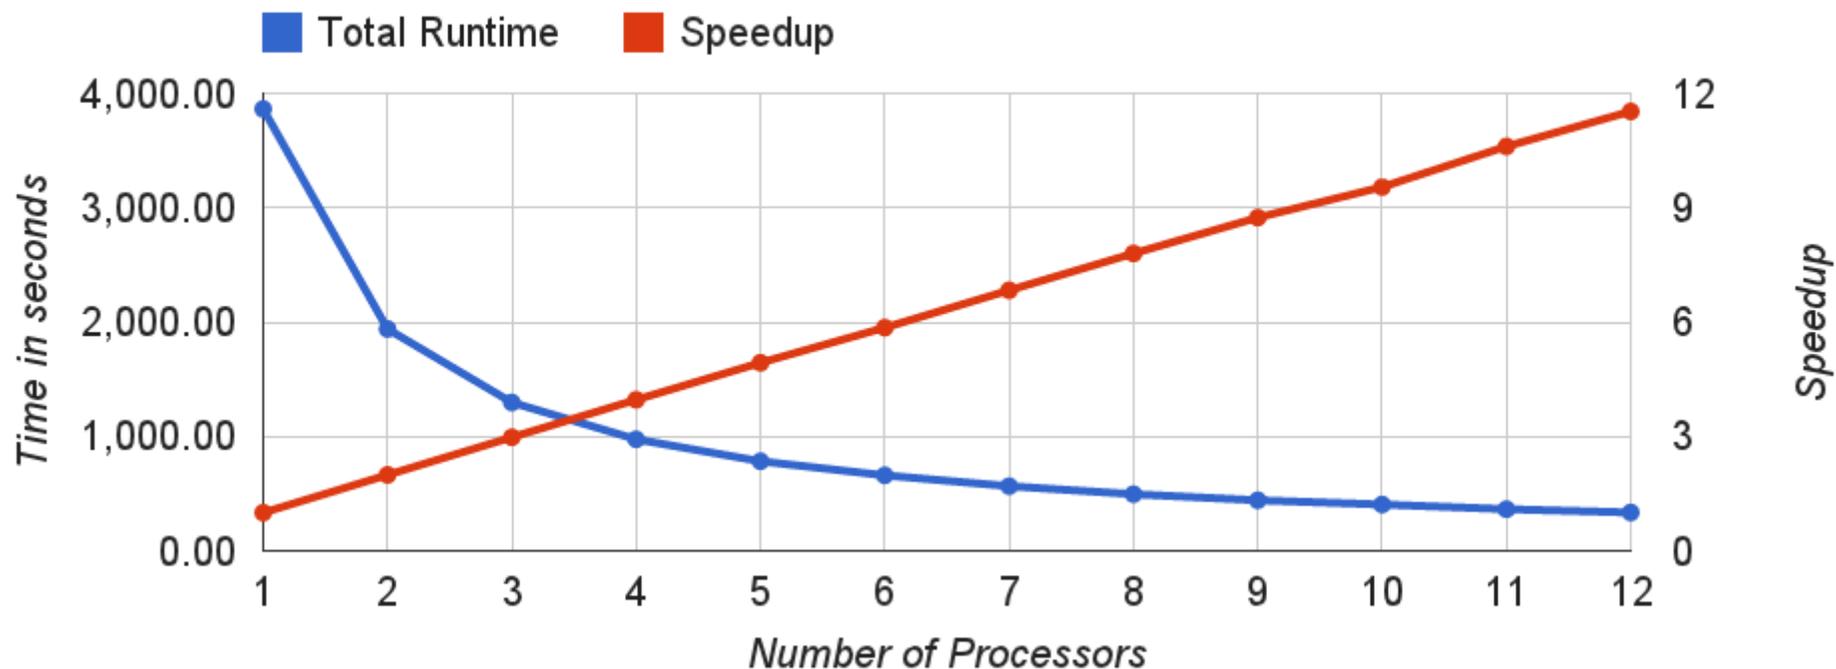


Benchmarks

Second Test: OpenMP on 12 Core System



- 12 Intel Xeon E5645 at 2.40GHz
- input: $\text{md5}(x)$ with $x \in \{0, 1\}^{32}$



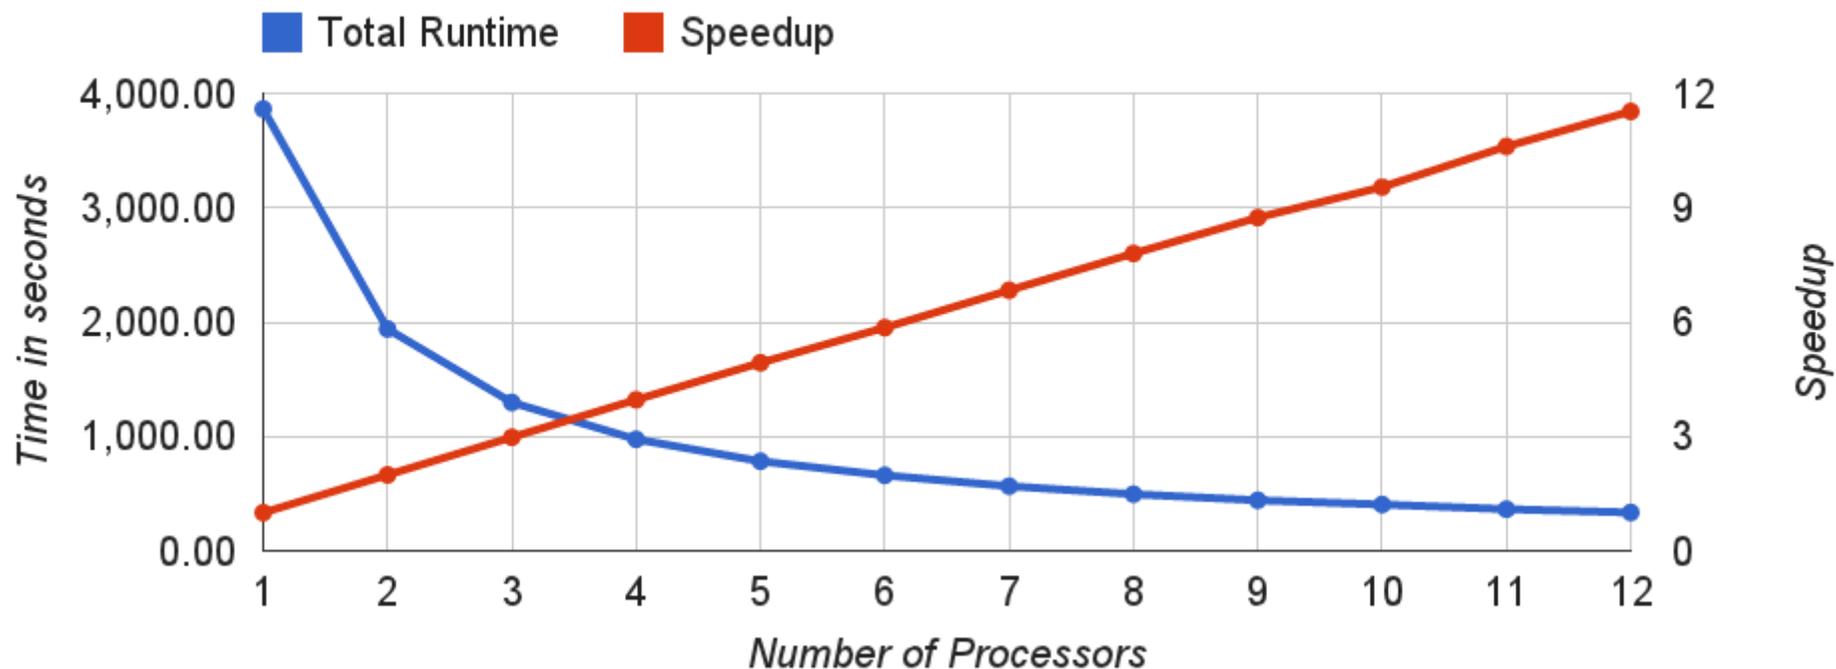
- 335.18 seconds to find collision

Benchmarks

Second Test: OpenMP on 12 Core System



- 12 Intel Xeon E5645 at 2.40GHz
- input: md5(x) with $x \in \{0, 1\}^{32}$



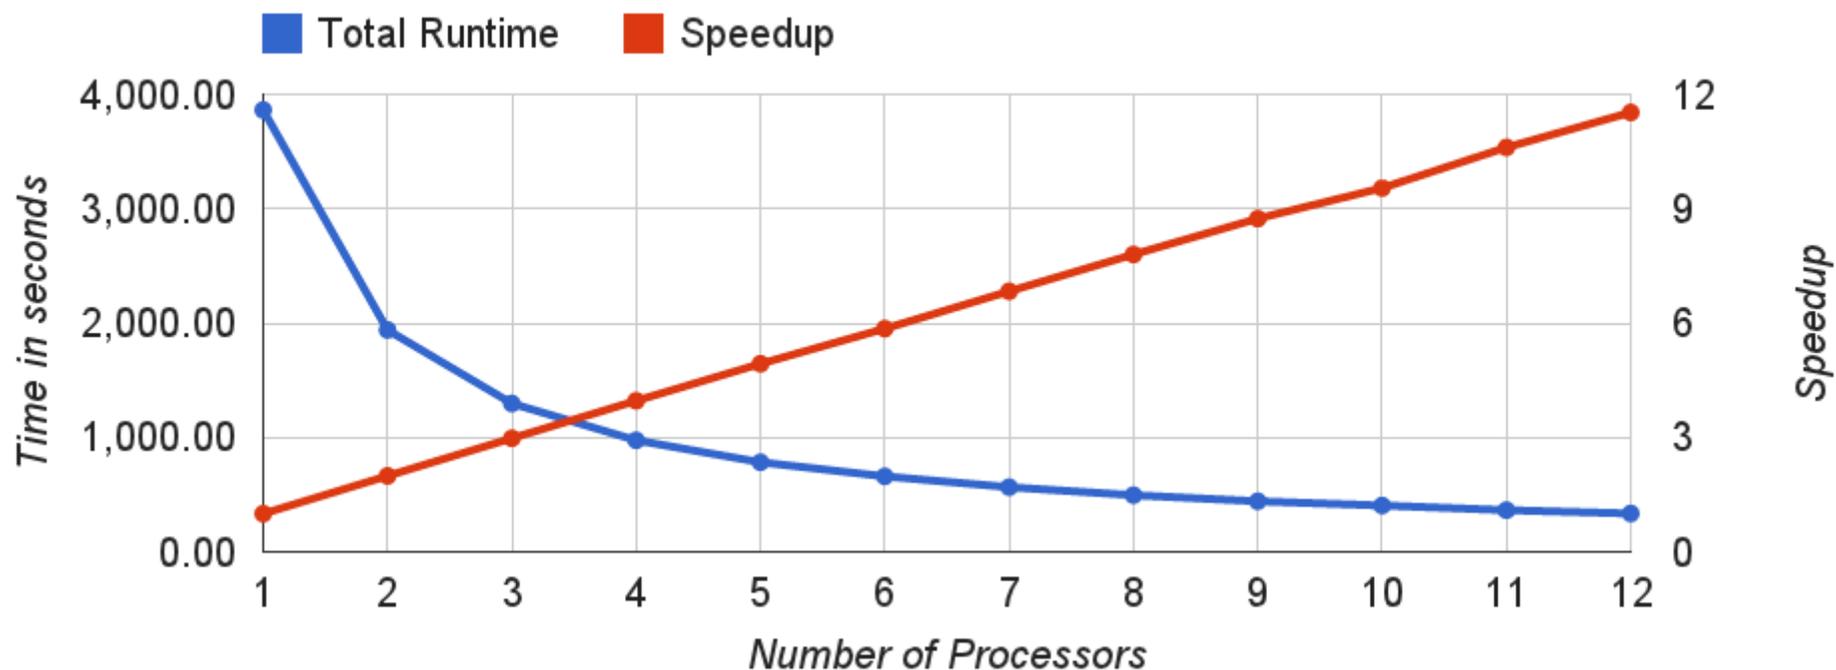
- 335.18 seconds to find collision
- approx. 1 billion strings tried

Benchmarks

Second Test: OpenMP on 12 Core System



- 12 Intel Xeon E5645 at 2.40GHz
- input: $\text{md5}(x)$ with $x \in \{0, 1\}^{32}$



- 335.18 seconds to find collision
- approx. 1 billion strings tried

→ linear speedup

Benchmarks



Benchmarks

Third Test: OpenMP on 32 Core System



Benchmarks

Third Test: OpenMP on 32 Core System



Benchmarks

Third Test: OpenMP on 32 Core System

- 32 Intel Xeon E7-4830 at 2.13GHz



Benchmarks

Third Test: OpenMP on 32 Core System

- 32 Intel Xeon E7-4830 at 2.13GHz
- input: $\text{md5}(x)$ with $x \in \{0, 1\}^{32}$

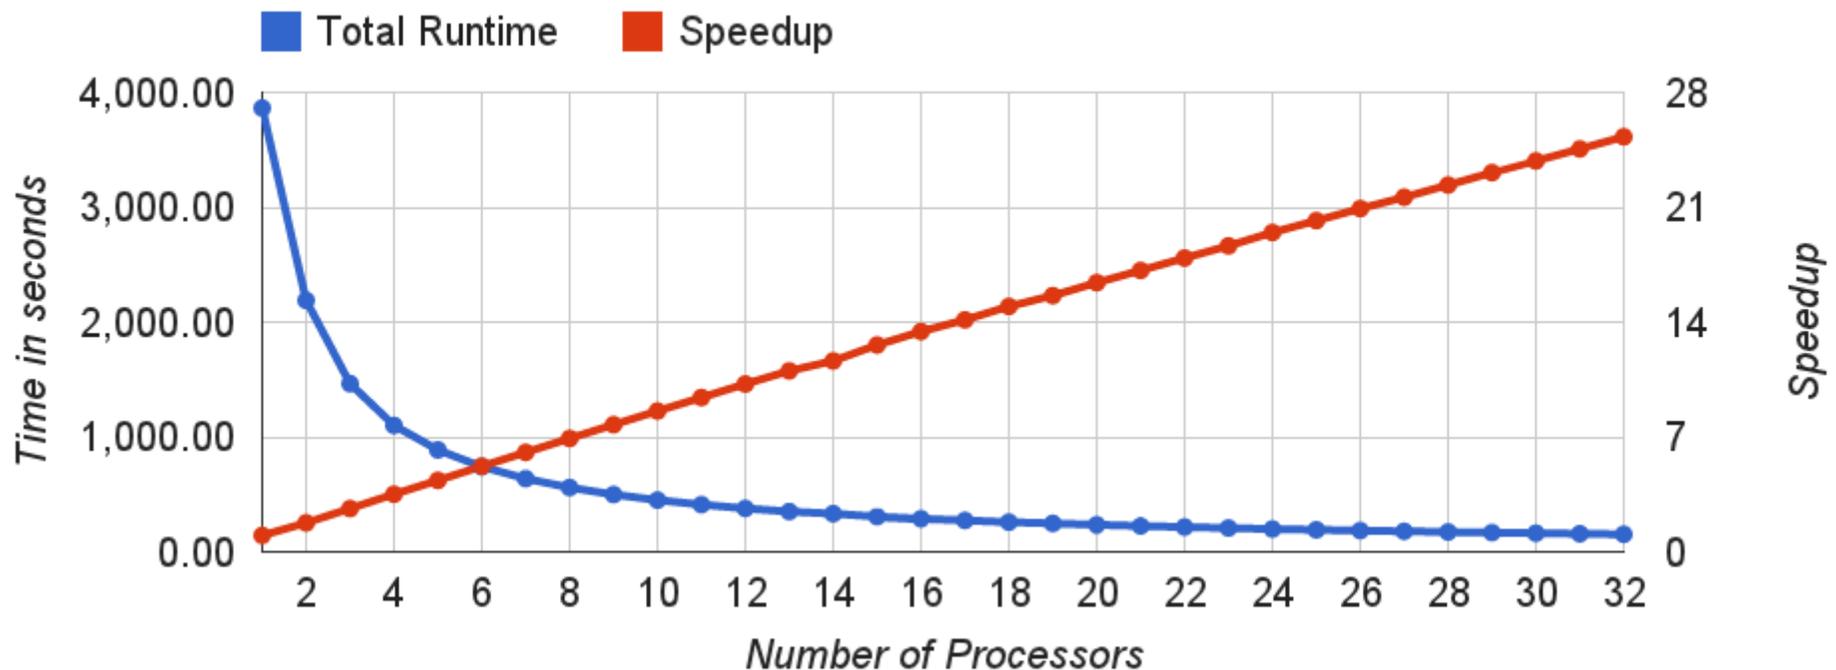


Benchmarks

Third Test: OpenMP on 32 Core System



- 32 Intel Xeon E7-4830 at 2.13GHz
- input: md5(x) with $x \in \{0, 1\}^{32}$

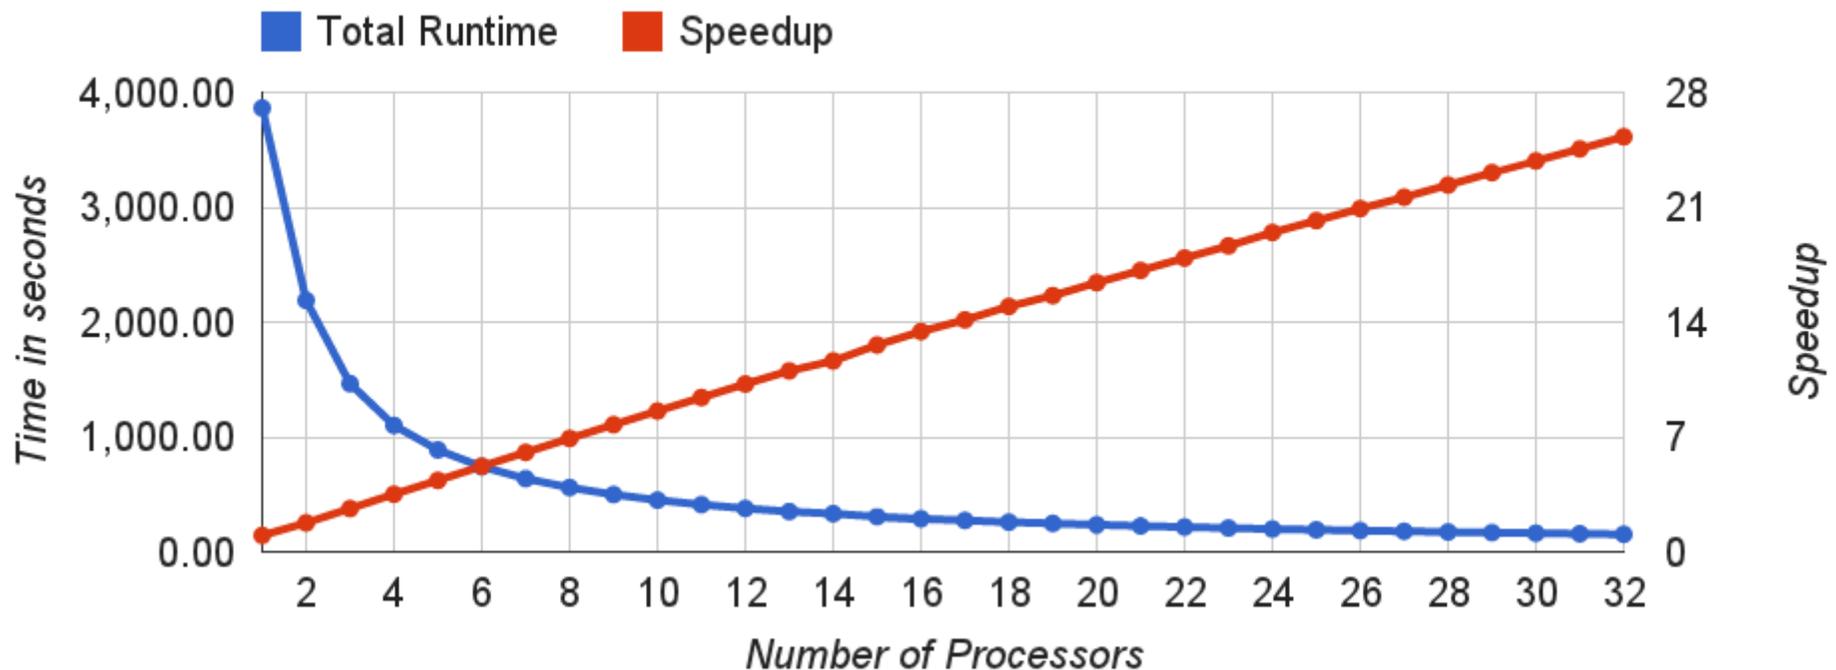


Benchmarks

Third Test: OpenMP on 32 Core System



- 32 Intel Xeon E7-4830 at 2.13GHz
- input: $\text{md5}(x)$ with $x \in \{0, 1\}^{32}$



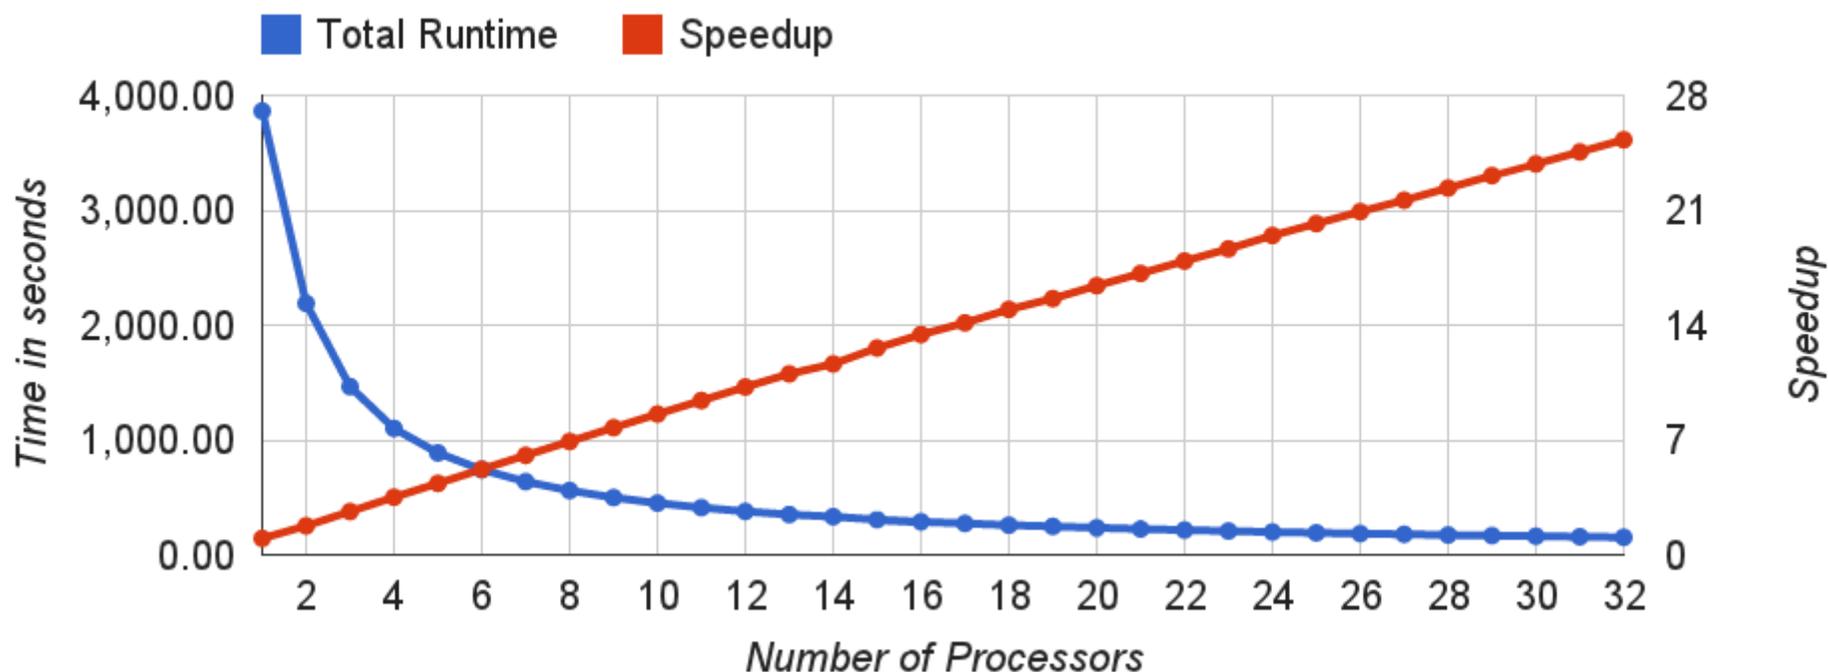
- 152.74 seconds to find collision

Benchmarks

Third Test: OpenMP on 32 Core System



- 32 Intel Xeon E7-4830 at 2.13GHz
- input: $\text{md5}(x)$ with $x \in \{0, 1\}^{32}$



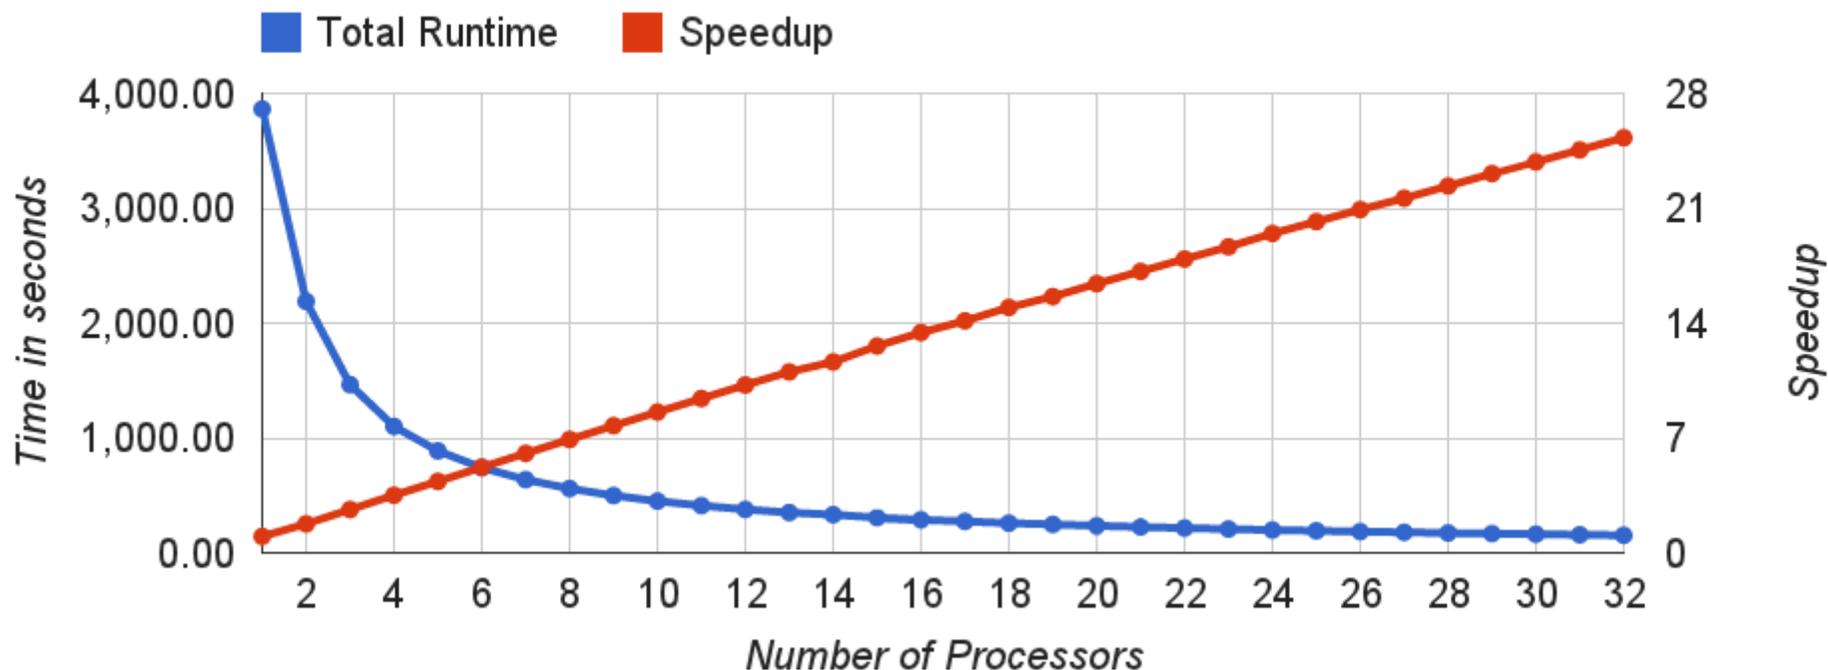
- 152.74 seconds to find collision
- approx. 1 billion strings tried

Benchmarks

Third Test: OpenMP on 32 Core System



- 32 Intel Xeon E7-4830 at 2.13GHz
- input: $\text{md5}(x)$ with $x \in \{0, 1\}^{32}$



- 152.74 seconds to find collision
- approx. 1 billion strings tried

→ linear speedup

Benchmarks



Benchmarks

Fourth Test: MPI on 6 · 12 Core Systems



Benchmarks

Fourth Test: MPI on 6 · 12 Core Systems



Benchmarks

Fourth Test: MPI on 6 · 12 Core Systems

- 6 hosts with 12 Intel Xeon E5645 at 2.40GHz



Benchmarks

Fourth Test: MPI on 6 · 12 Core Systems



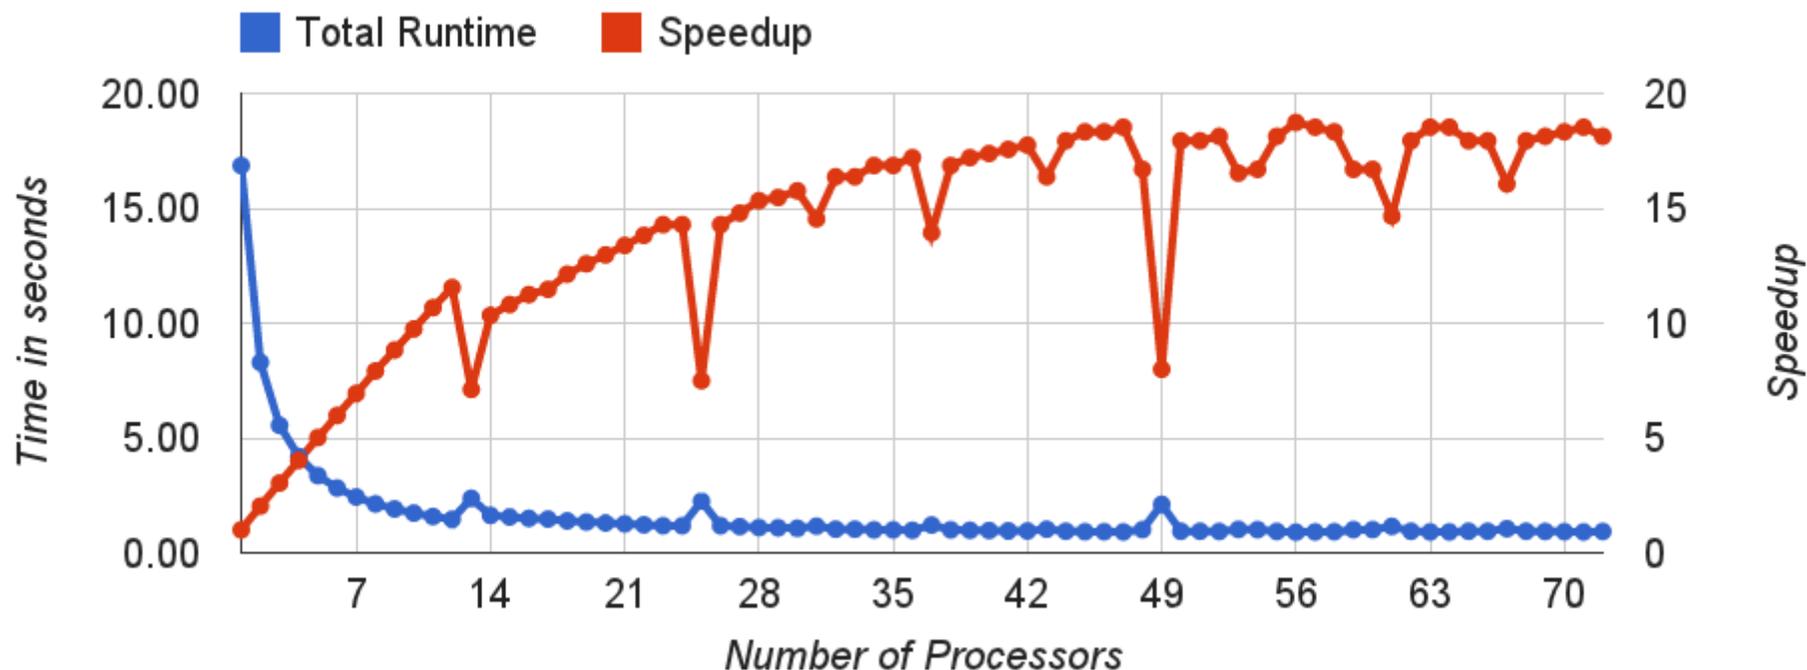
- 6 hosts with 12 Intel Xeon E5645 at 2.40GHz
- input: $\text{md5}(x)$ with $x \in \{0, 1\}^{24}$

Benchmarks

Fourth Test: MPI on 6 · 12 Core Systems



- 6 hosts with 12 Intel Xeon E5645 at 2.40GHz
- input: md5(x) with $x \in \{0, 1\}^{24}$

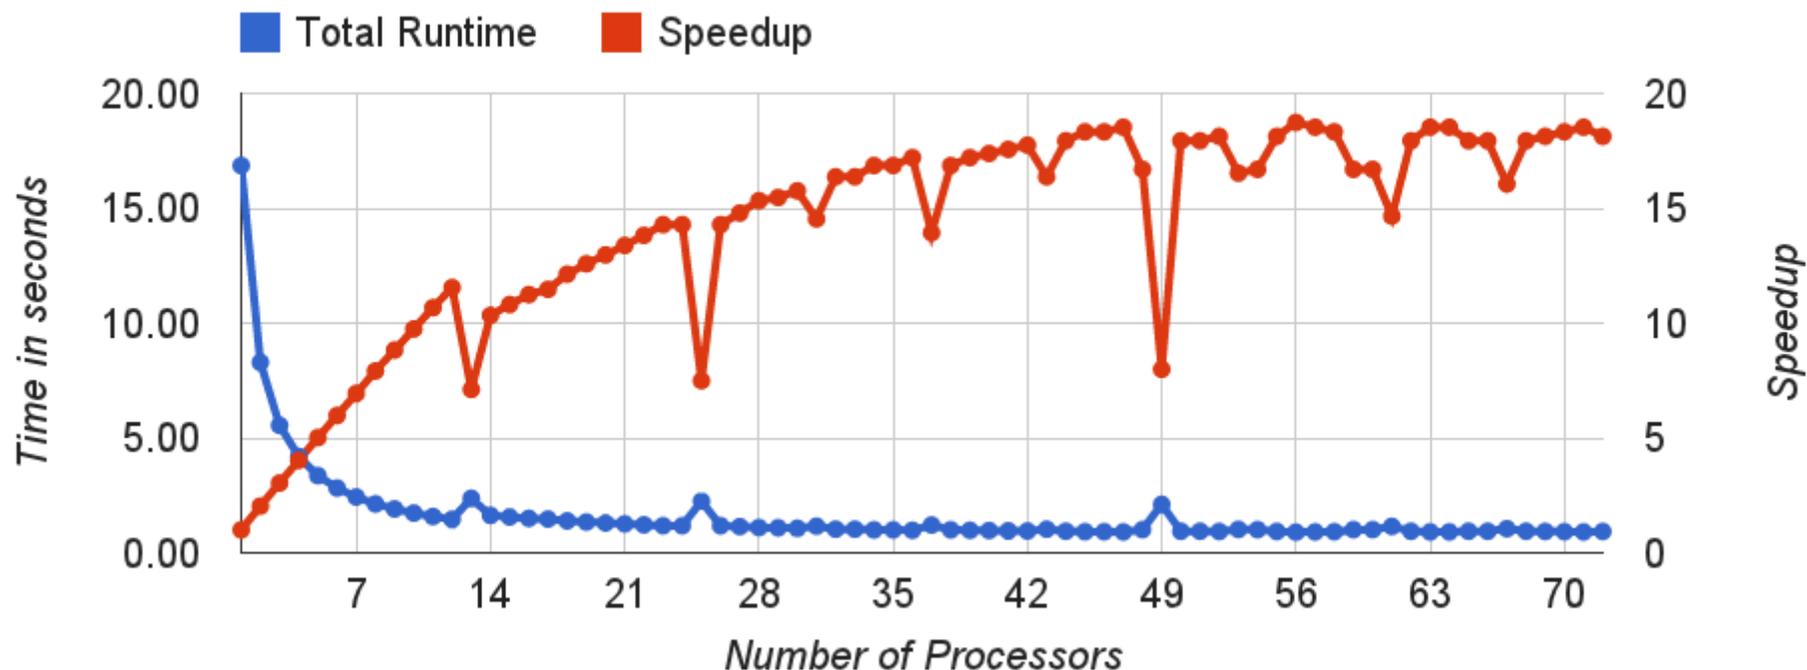


Benchmarks

Fourth Test: MPI on 6 · 12 Core Systems



- 6 hosts with 12 Intel Xeon E5645 at 2.40GHz
- input: md5(x) with $x \in \{0, 1\}^{24}$



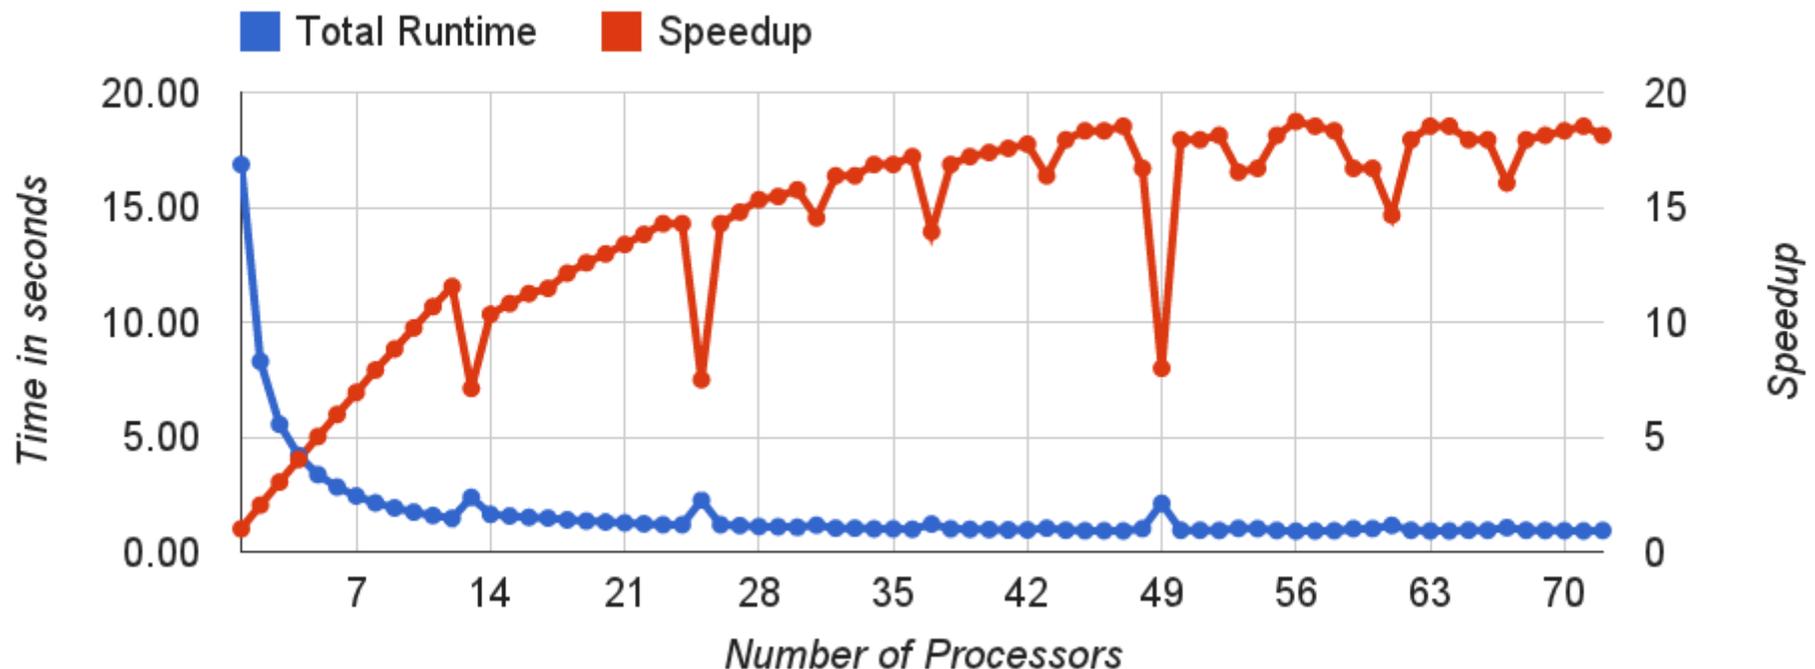
- 0.93 seconds to find collision

Benchmarks

Fourth Test: MPI on 6 · 12 Core Systems



- 6 hosts with 12 Intel Xeon E5645 at 2.40GHz
- input: md5(x) with $x \in \{0, 1\}^{24}$



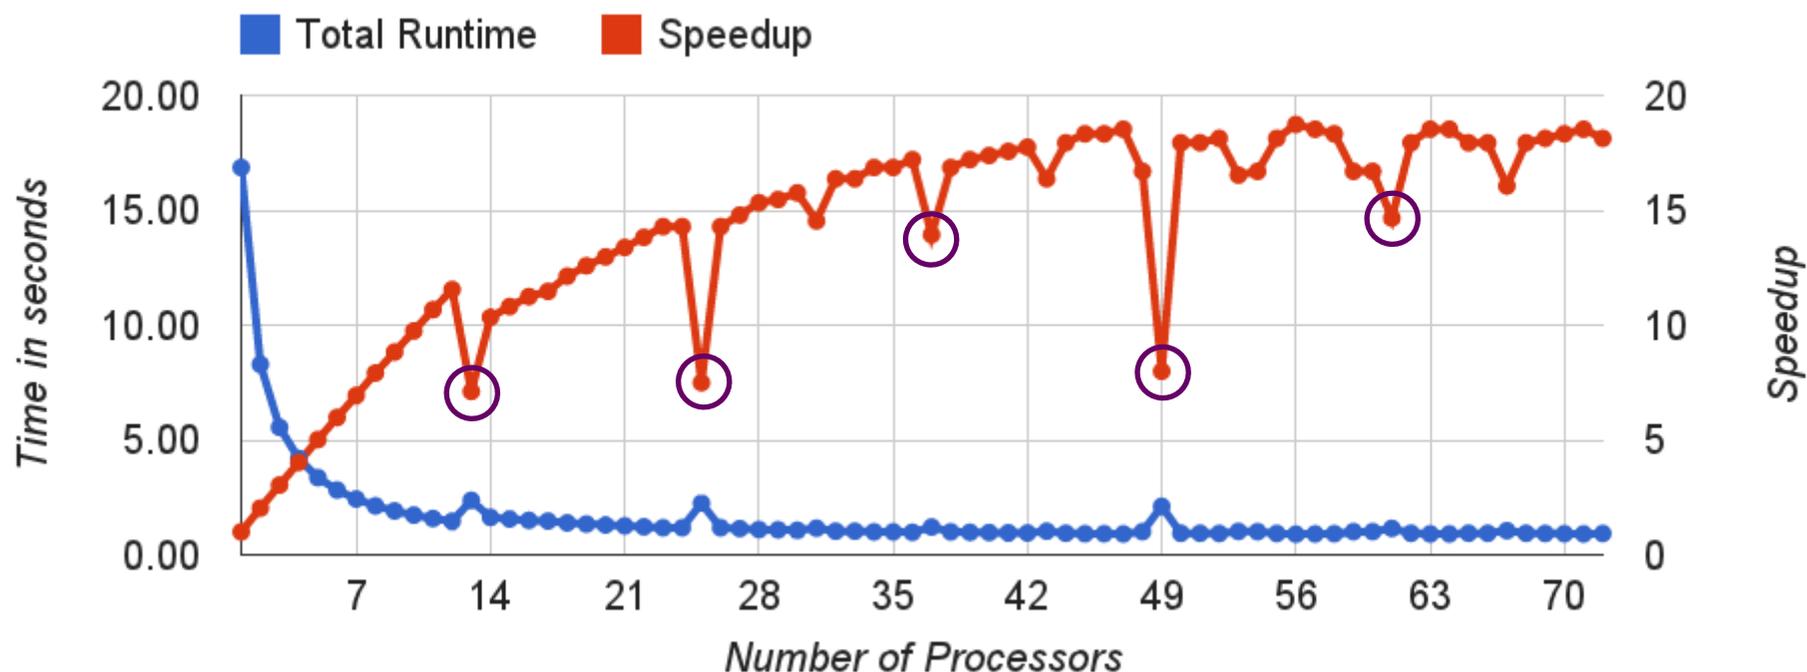
- 0.93 seconds to find collision
- approx. 4 million strings tried

Benchmarks

Fourth Test: MPI on 6 · 12 Core Systems



- 6 hosts with 12 Intel Xeon E5645 at 2.40GHz
- input: $\text{md5}(x)$ with $x \in \{0, 1\}^{24}$



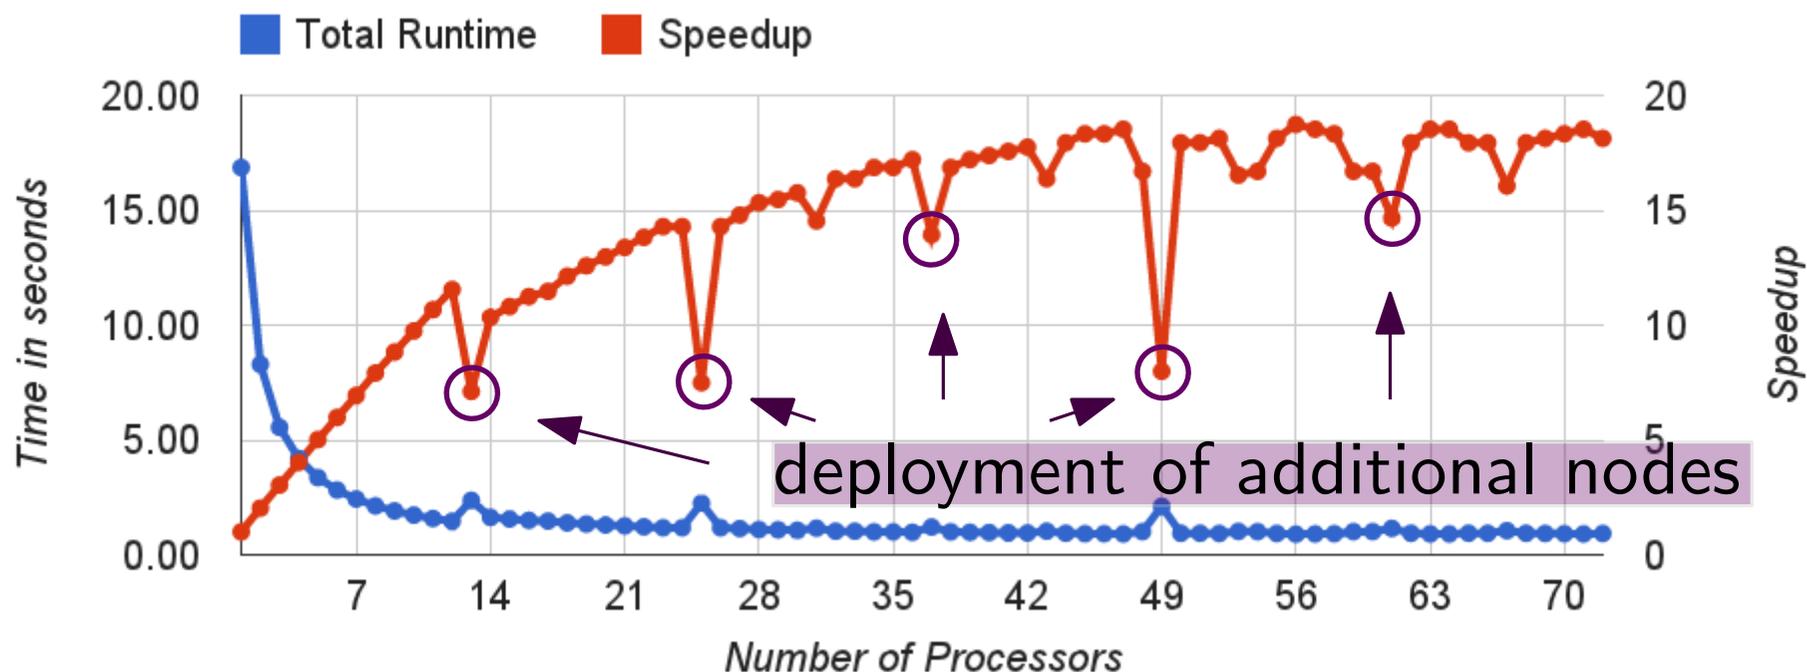
- 0.93 seconds to find collision → speed drop
- approx. 4 million strings tried

Benchmarks

Fourth Test: MPI on 6 · 12 Core Systems



- 6 hosts with 12 Intel Xeon E5645 at 2.40GHz
- input: $\text{md5}(x)$ with $x \in \{0, 1\}^{24}$



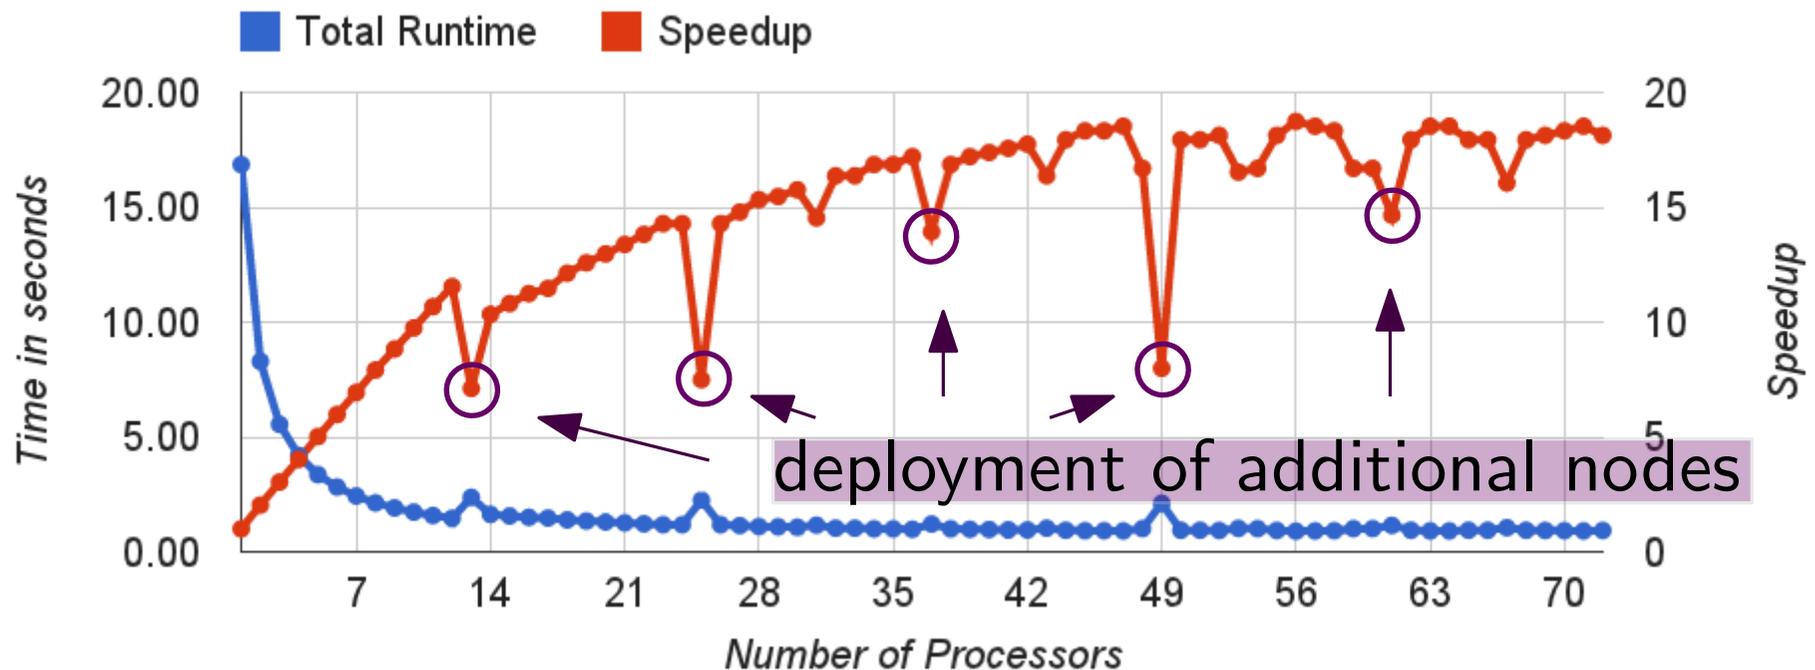
- 0.93 seconds to find collision \longrightarrow speed drop
- approx. 4 million strings tried

Benchmarks

Fourth Test: MPI on 6 · 12 Core Systems



- 6 hosts with 12 Intel Xeon E5645 at 2.40GHz
- input: $\text{md5}(x)$ with $x \in \{0, 1\}^{24}$



- 0.93 seconds to find collision
 - approx. 4 million strings tried
- speed drop due to
- small problem size
 - setup overhead

Benchmarks



Benchmarks

Fifth Test: MPI on 6 · 12 Core Systems



Benchmarks

Fifth Test: MPI on 6 · 12 Core Systems

- 6 hosts with 12 Intel Xeon E5645 at 2.40GHz



Benchmarks

Fifth Test: MPI on 6 · 12 Core Systems



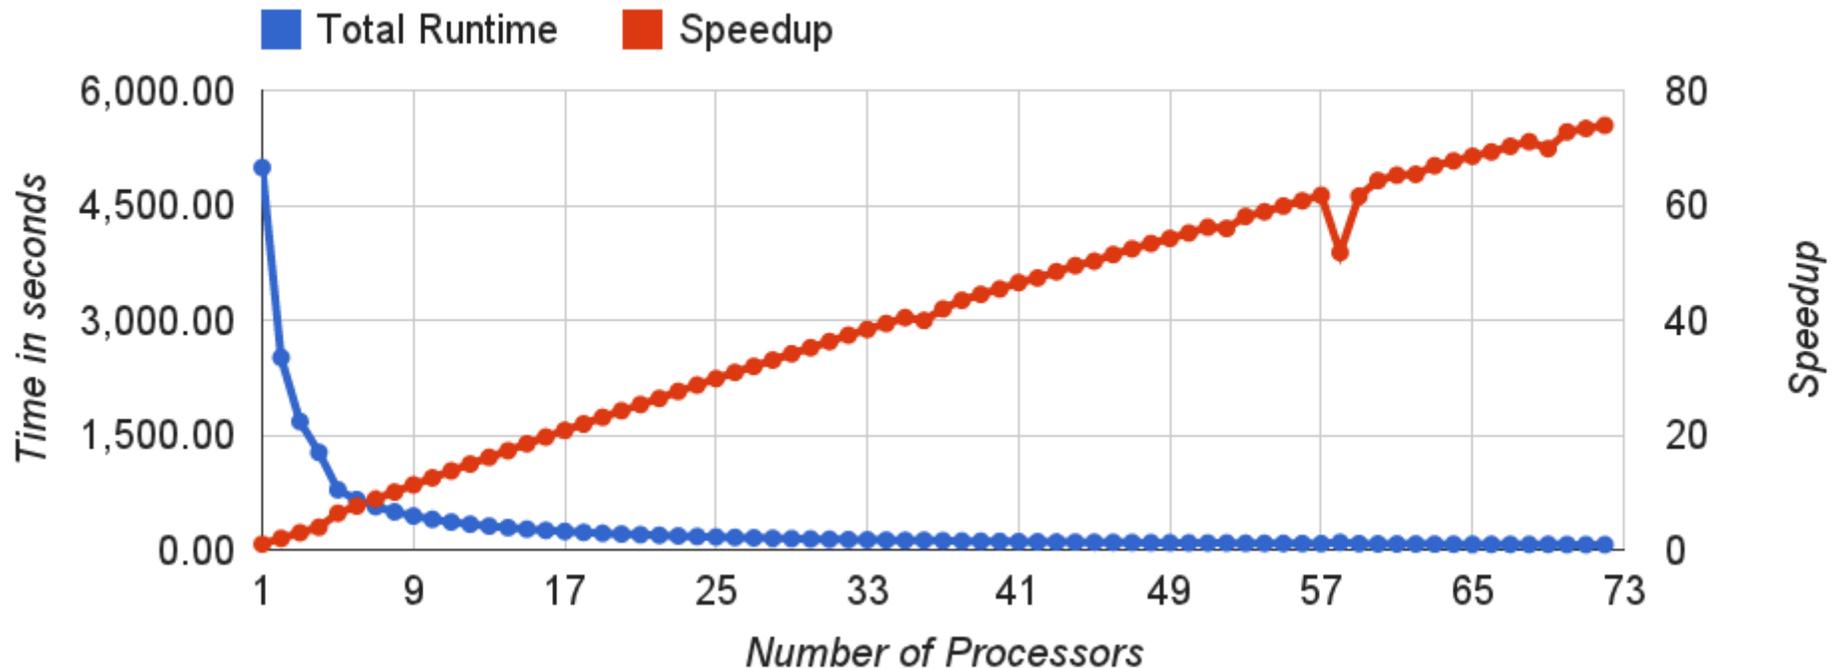
- 6 hosts with 12 Intel Xeon E5645 at 2.40GHz
- input: $\text{md5}(x)$ with $x \in \{0, 1\}^{32}$

Benchmarks

Fifth Test: MPI on 6 · 12 Core Systems



- 6 hosts with 12 Intel Xeon E5645 at 2.40GHz
- input: md5(x) with $x \in \{0, 1\}^{32}$

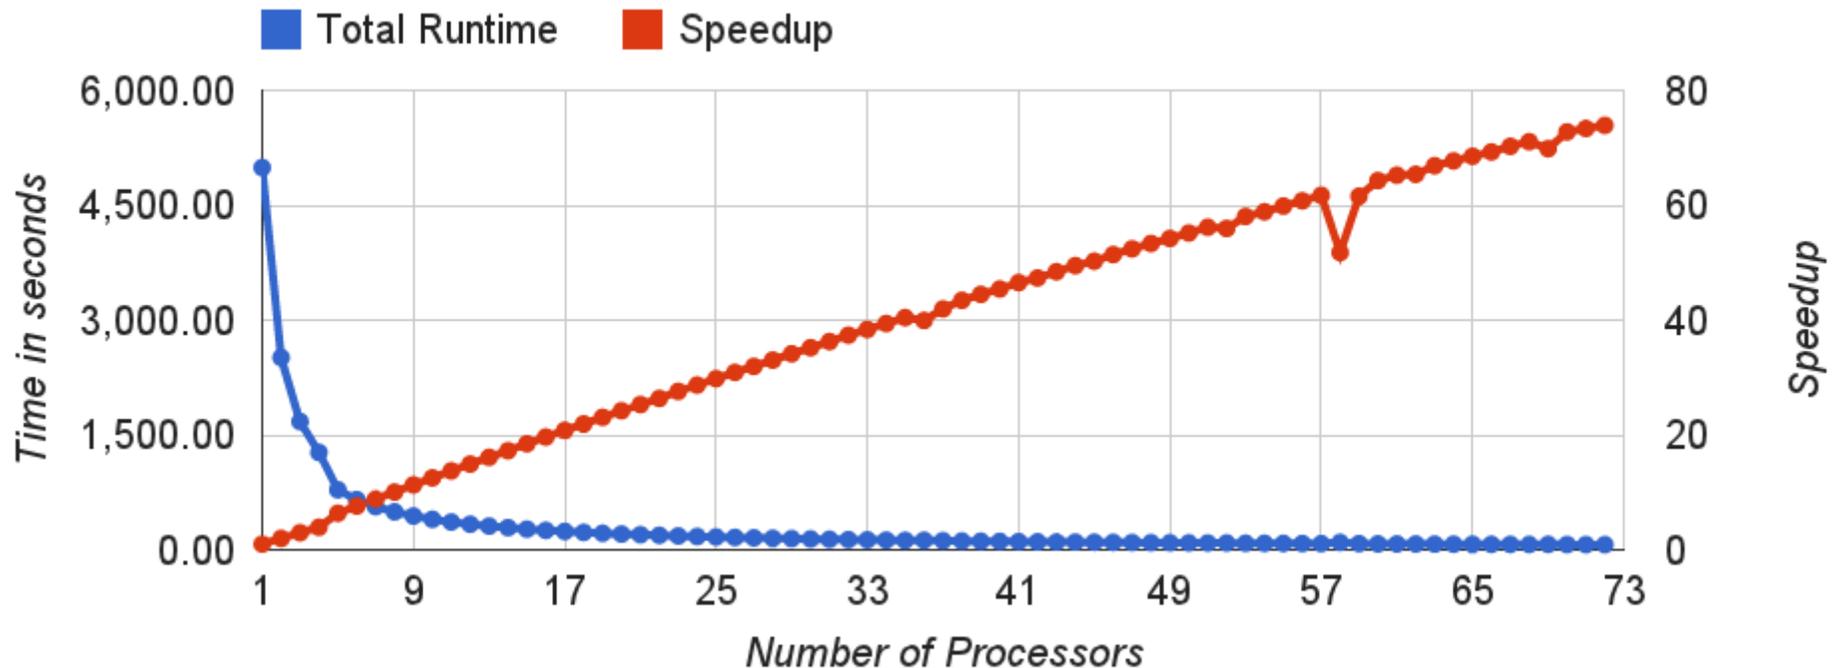


Benchmarks

Fifth Test: MPI on 6 · 12 Core Systems



- 6 hosts with 12 Intel Xeon E5645 at 2.40GHz
- input: md5(x) with $x \in \{0, 1\}^{32}$



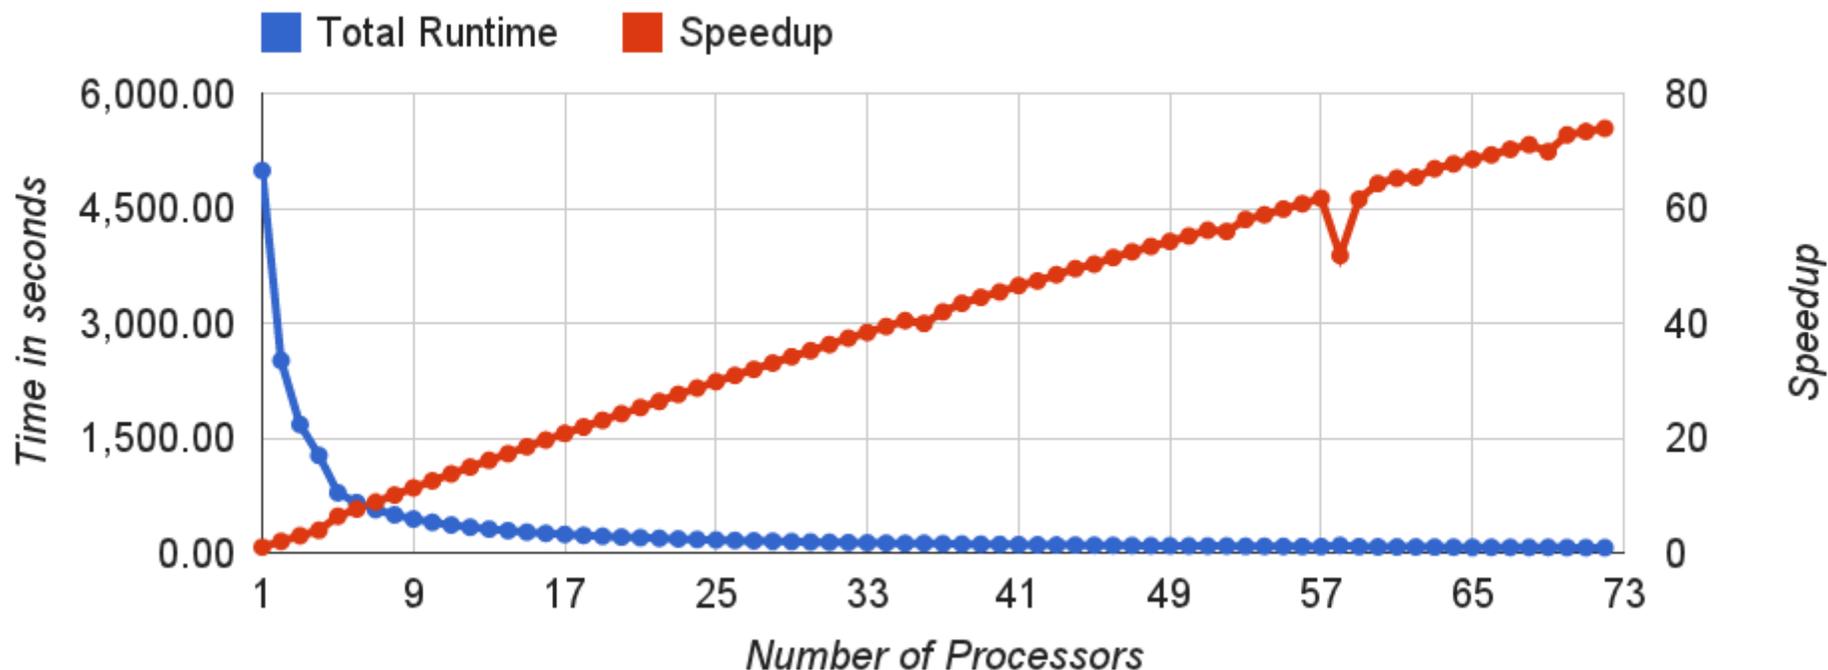
- 67.54 seconds to find collision

Benchmarks

Fifth Test: MPI on 6 · 12 Core Systems



- 6 hosts with 12 Intel Xeon E5645 at 2.40GHz
- input: md5(x) with $x \in \{0, 1\}^{32}$



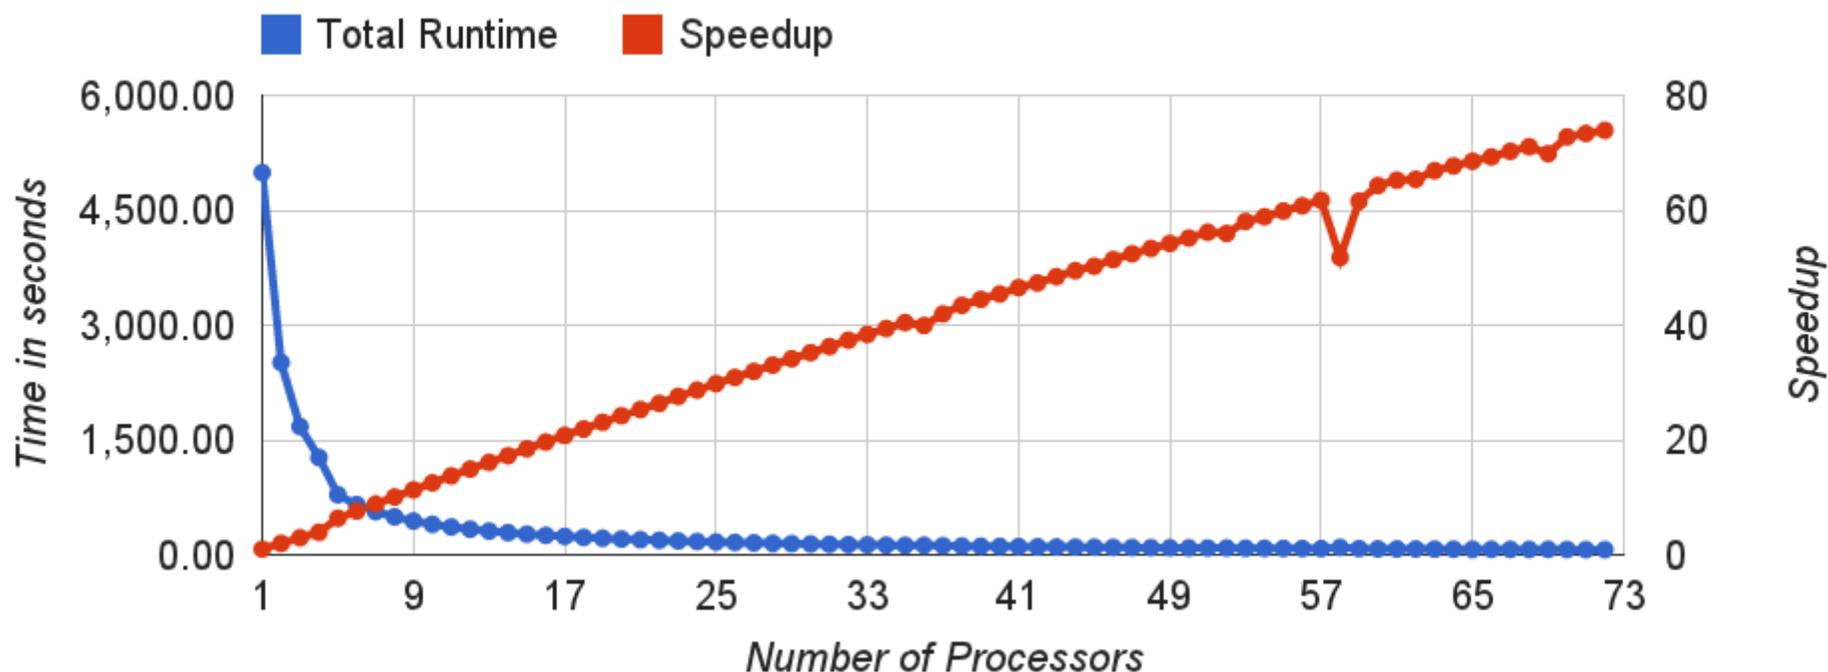
- 67.54 seconds to find collision
- approx. 1 billion strings tried

Benchmarks



Fifth Test: MPI on 6 · 12 Core Systems

- 6 hosts with 12 Intel Xeon E5645 at 2.40GHz
- input: md5(x) with $x \in \{0, 1\}^{32}$



- 67.54 seconds to find collision
- approx. 1 billion strings tried

→ linear speedup

Benchmarks



Benchmarks

Sixth Test: MPI on 2 · 32 Core Systems



Benchmarks

Sixth Test: MPI on 2 · 32 Core Systems



Benchmarks

Sixth Test: MPI on 2 · 32 Core Systems

- 2 hosts with 32 Xeon E7-4830 at 2.13GHz



Benchmarks

Sixth Test: MPI on 2 · 32 Core Systems

- 2 hosts with 32 Xeon E7-4830 at 2.13GHz
- input: $\text{md5}(x)$ with $x \in \{0, 1\}^{24}$

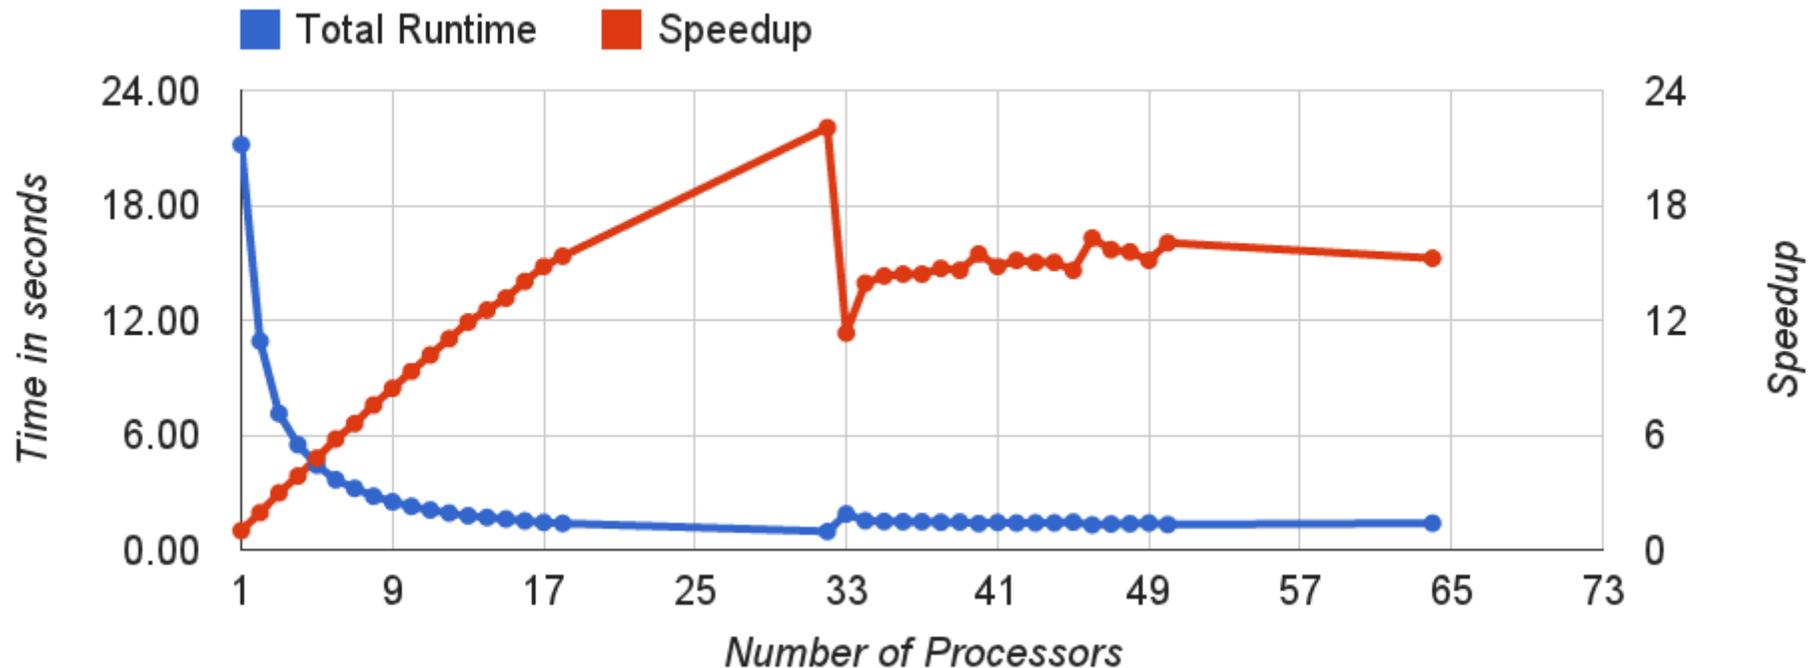


Benchmarks

Sixth Test: MPI on 2 · 32 Core Systems



- 2 hosts with 32 Xeon E7-4830 at 2.13GHz
- input: md5(x) with $x \in \{0, 1\}^{24}$

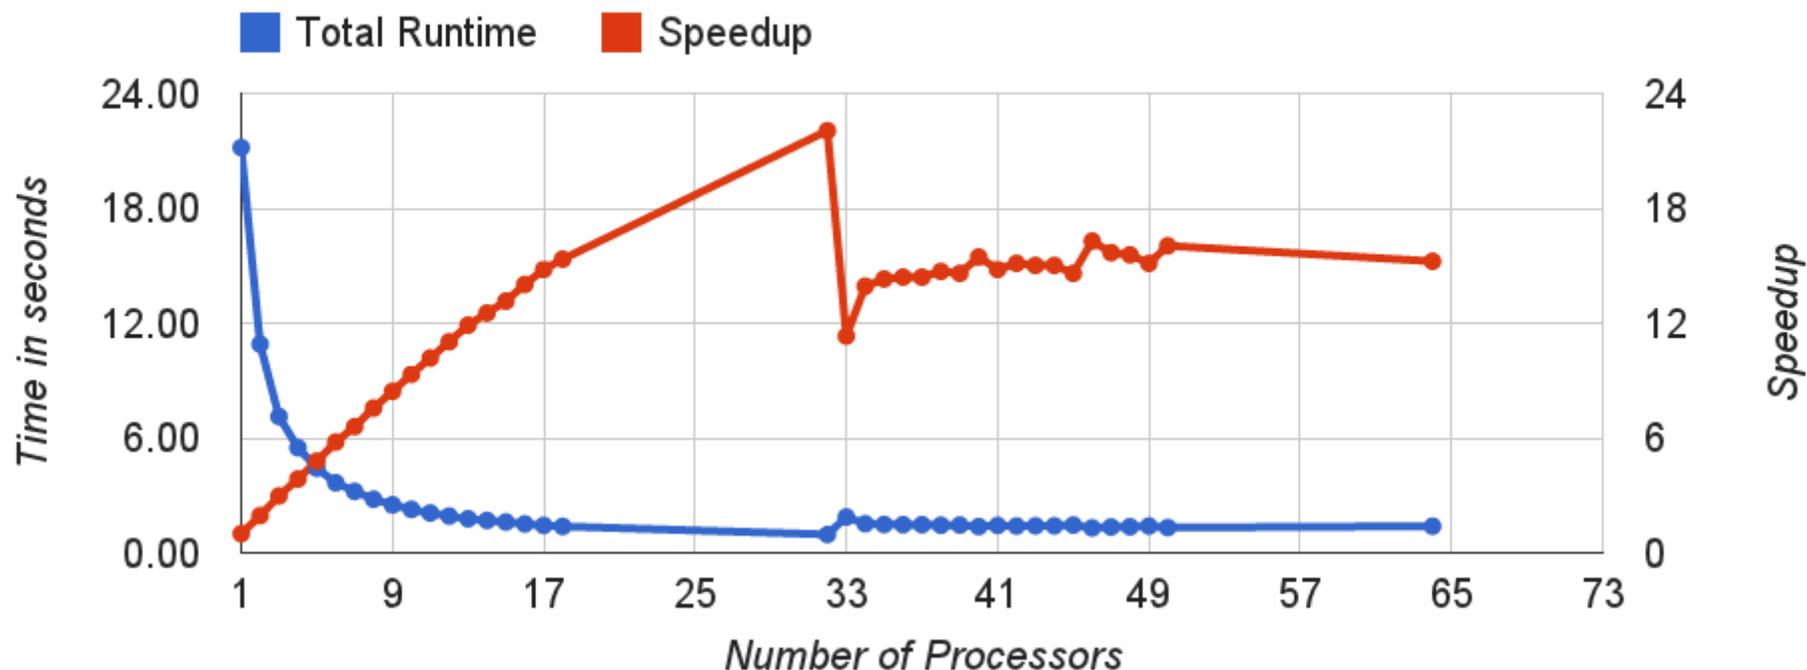


Benchmarks

Sixth Test: MPI on 2 · 32 Core Systems



- 2 hosts with 32 Xeon E7-4830 at 2.13GHz
- input: $\text{md5}(x)$ with $x \in \{0, 1\}^{24}$



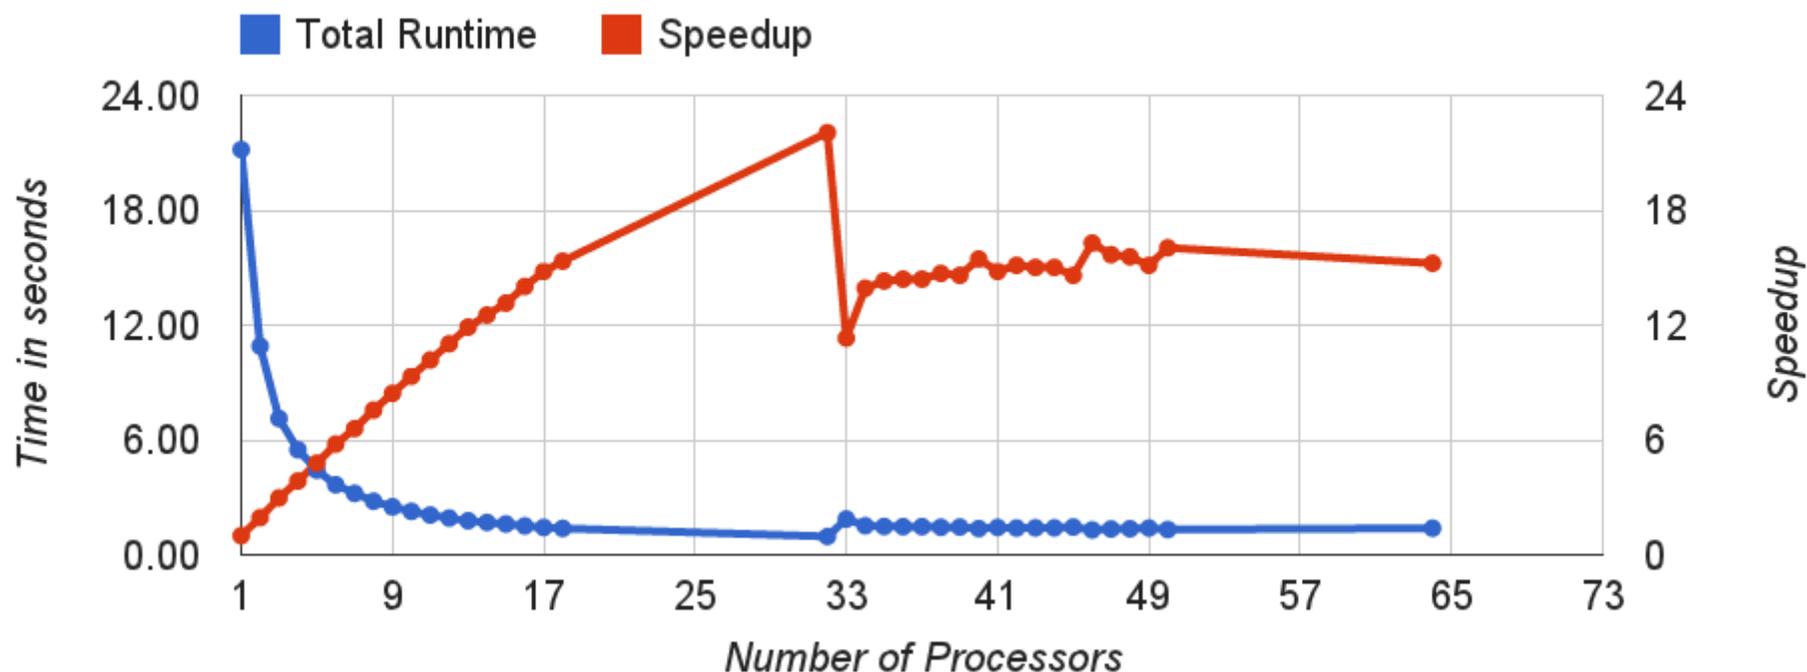
- 0.96 seconds to find collision

Benchmarks

Sixth Test: MPI on 2 · 32 Core Systems



- 2 hosts with 32 Xeon E7-4830 at 2.13GHz
- input: $\text{md5}(x)$ with $x \in \{0, 1\}^{24}$



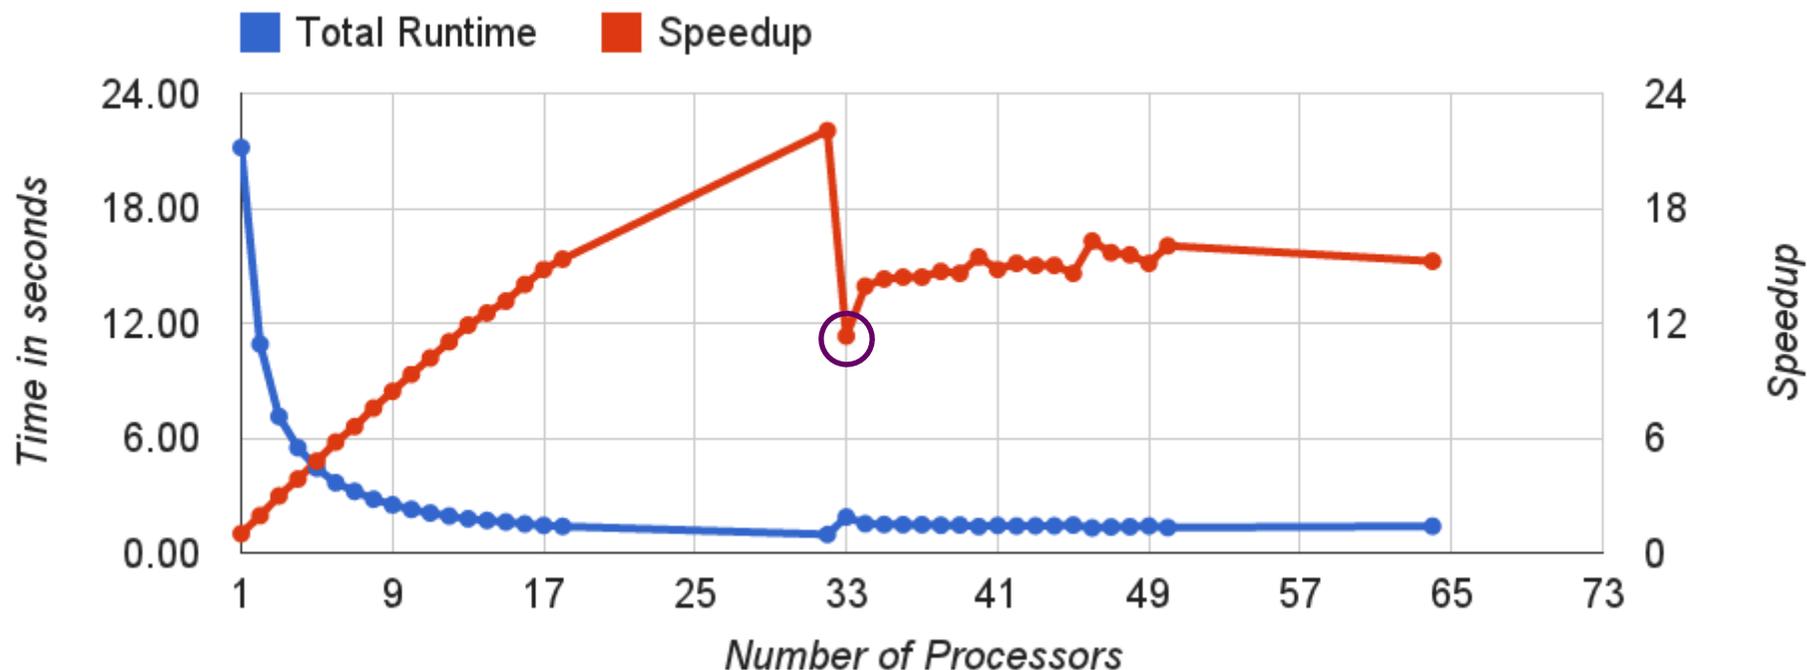
- 0.96 seconds to find collision
- approx. 4 million strings tried

Benchmarks



Sixth Test: MPI on 2 · 32 Core Systems

- 2 hosts with 32 Xeon E7-4830 at 2.13GHz
- input: $\text{md5}(x)$ with $x \in \{0, 1\}^{24}$



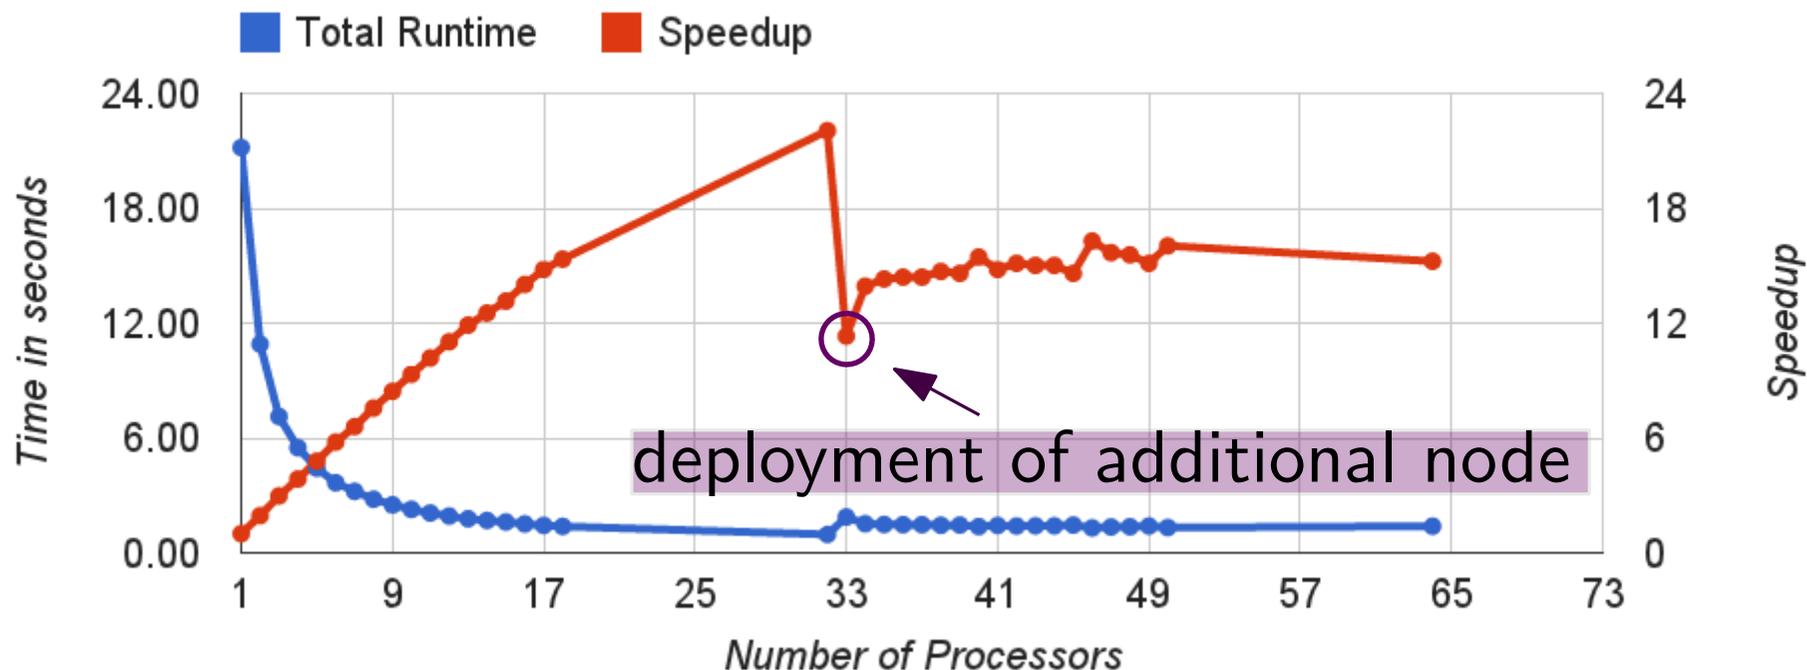
- 0.96 seconds to find collision \longrightarrow speed drop
- approx. 4 million strings tried

Benchmarks



Sixth Test: MPI on 2 · 32 Core Systems

- 2 hosts with 32 Xeon E7-4830 at 2.13GHz
- input: $\text{md5}(x)$ with $x \in \{0, 1\}^{24}$



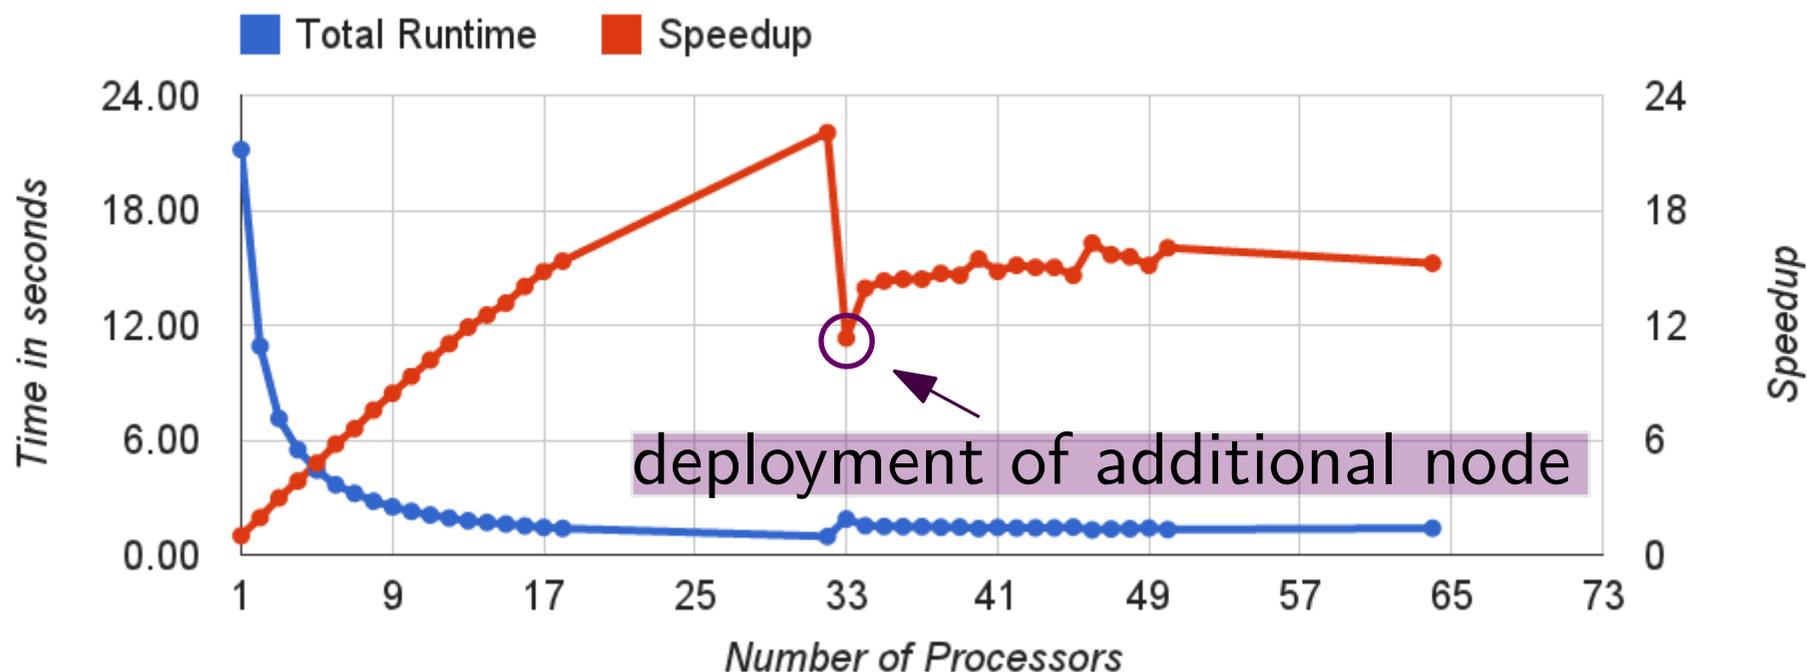
- 0.96 seconds to find collision \longrightarrow speed drop due to
- approx. 4 million strings tried

Benchmarks



Sixth Test: MPI on 2 · 32 Core Systems

- 2 hosts with 32 Xeon E7-4830 at 2.13GHz
- input: $\text{md5}(x)$ with $x \in \{0, 1\}^{24}$



- 0.96 seconds to find collision
 - approx. 4 million strings tried
- ▶ speed drop due to
- small problem size
 - setup overhead

Benchmarks



Benchmarks

Seventh Test: MPI on 2 · 32 Core Systems



Benchmarks

Seventh Test: MPI on 2 · 32 Core Systems



- 2 hosts with 32 Xeon E7-4830 at 2.13GHz

Benchmarks

Seventh Test: MPI on 2 · 32 Core Systems



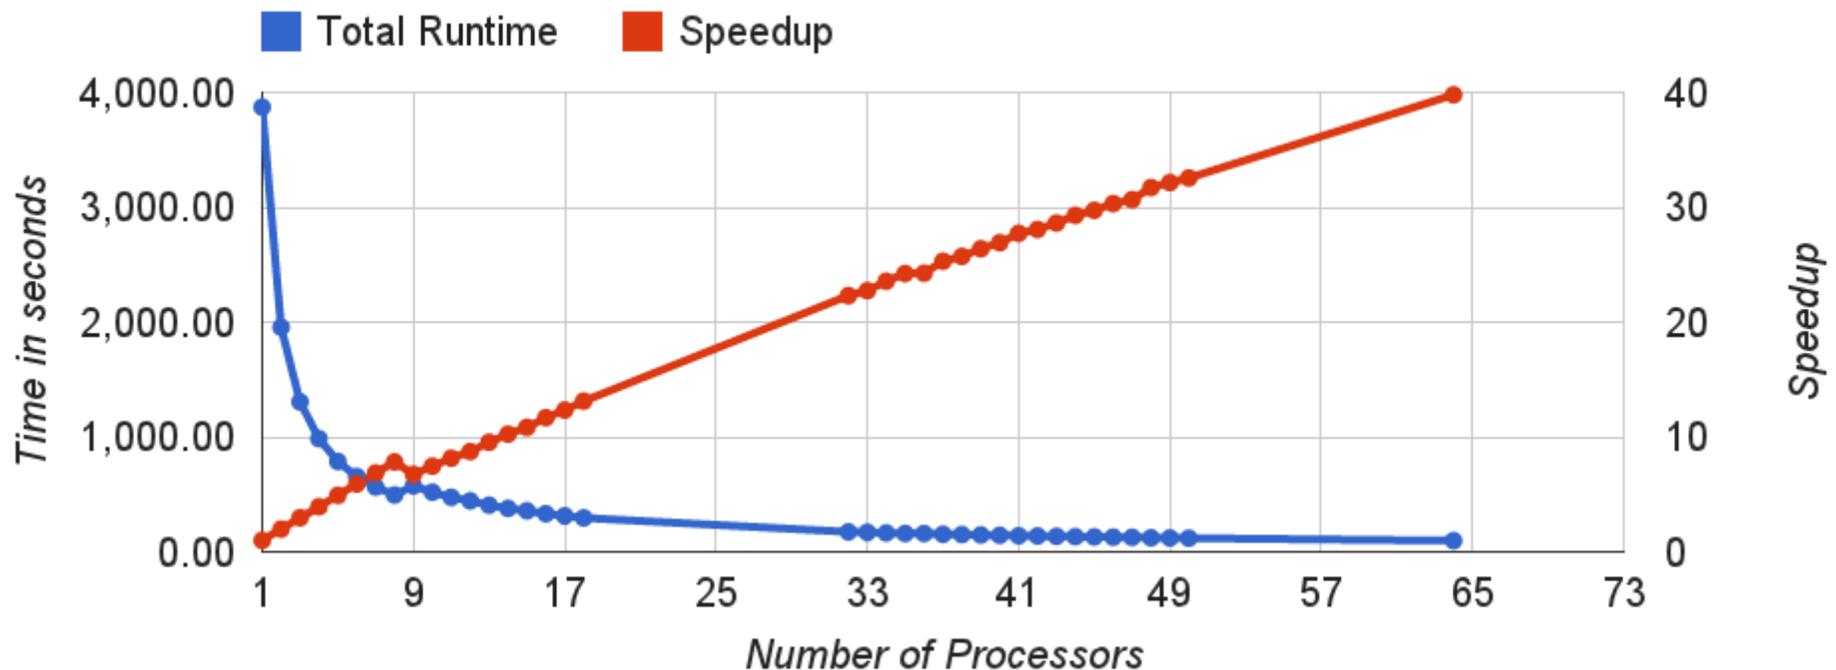
- 2 hosts with 32 Xeon E7-4830 at 2.13GHz
- input: $\text{md5}(x)$ with $x \in \{0, 1\}^{32}$

Benchmarks

Seventh Test: MPI on 2 · 32 Core Systems



- 2 hosts with 32 Xeon E7-4830 at 2.13GHz
- input: md5(x) with $x \in \{0, 1\}^{32}$

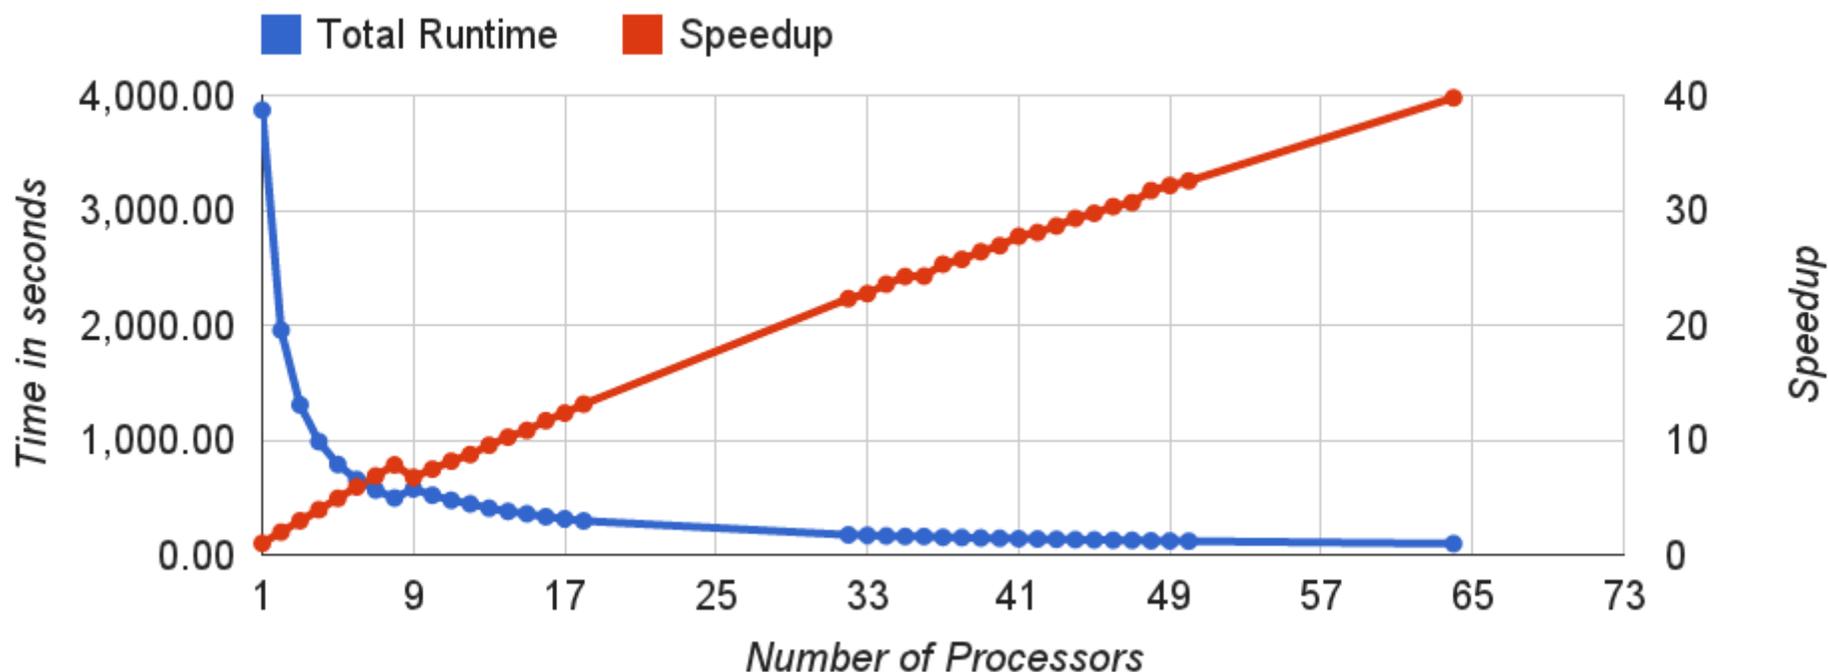


Benchmarks

Seventh Test: MPI on 2 · 32 Core Systems



- 2 hosts with 32 Xeon E7-4830 at 2.13GHz
- input: $\text{md5}(x)$ with $x \in \{0, 1\}^{32}$



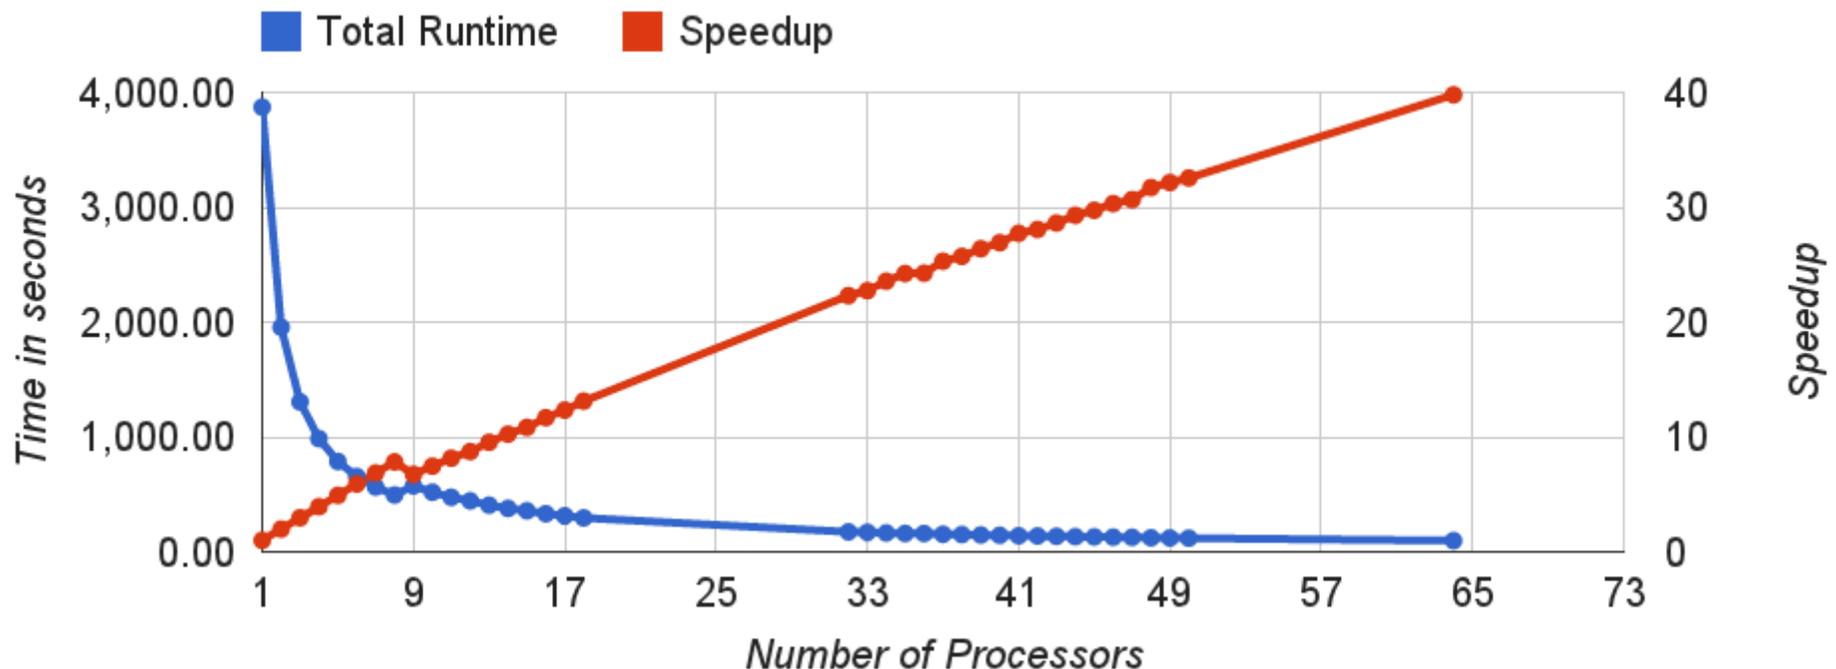
- 97.30 seconds to find collision

Benchmarks

Seventh Test: MPI on 2 · 32 Core Systems



- 2 hosts with 32 Xeon E7-4830 at 2.13GHz
- input: $\text{md5}(x)$ with $x \in \{0, 1\}^{32}$



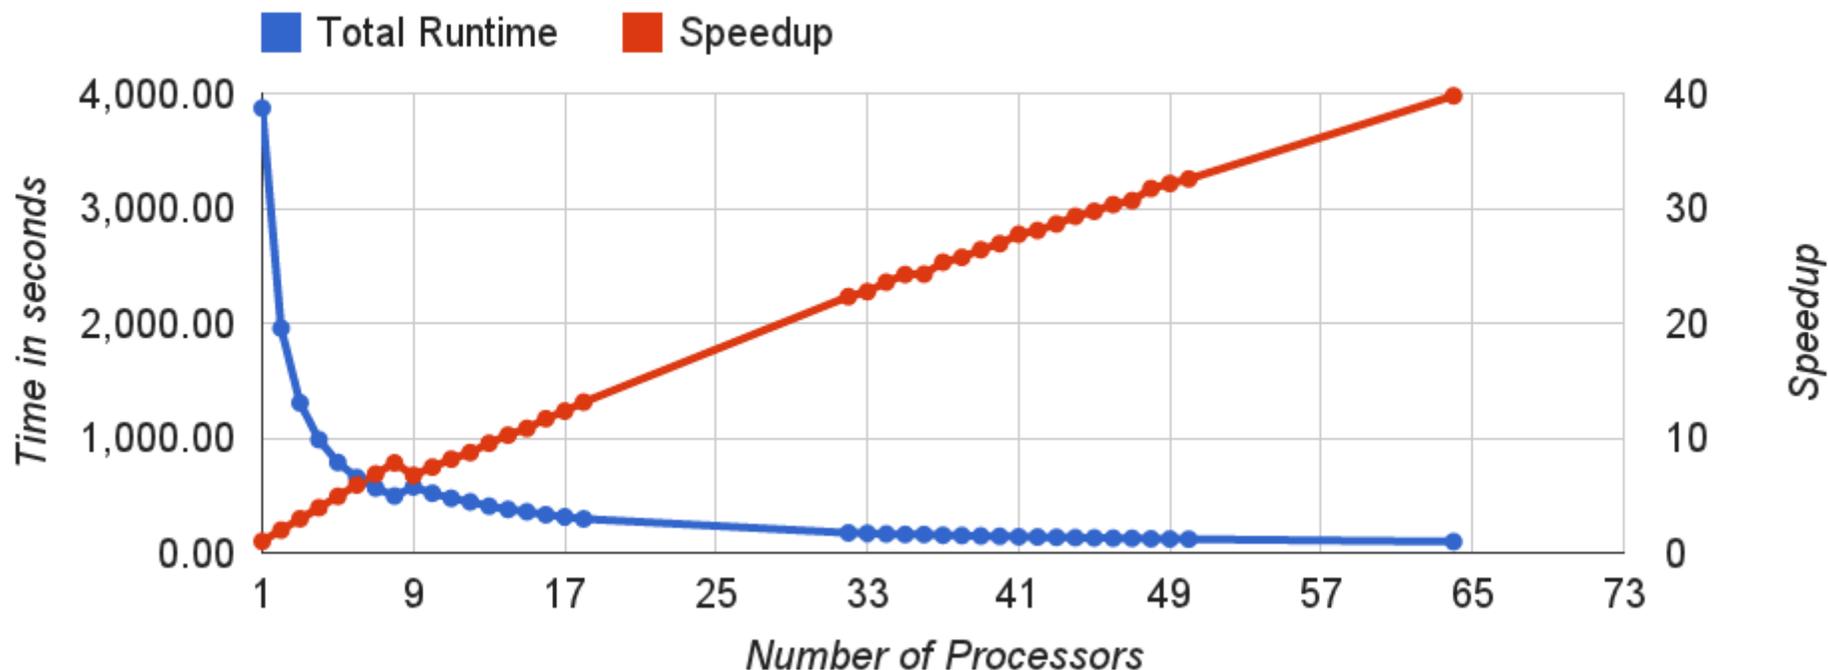
- 97.30 seconds to find collision
- approx. 1 billion strings tried

Benchmarks

Seventh Test: MPI on 2 · 32 Core Systems



- 2 hosts with 32 Xeon E7-4830 at 2.13GHz
- input: $\text{md5}(x)$ with $x \in \{0, 1\}^{32}$



- 97.30 seconds to find collision
- approx. 1 billion strings tried

→ linear speedup

Explanations and Expectations



Explanations and Expectations

OpenMP implementation



Explanations and Expectations

OpenMP implementation

- linear speedup due to few communication



Explanations and Expectations

OpenMP implementation

- linear speedup due to few communication
- suitable for smaller problem sizes



Explanations and Expectations

OpenMP implementation

- linear speedup due to few communication
- suitable for smaller problem sizes
 - but: already used CCR's "biggest" machine



Explanations and Expectations



OpenMP implementation

- linear speedup due to few communication
- suitable for smaller problem sizes
 - but: already used CCR's "biggest" machine

MPI implementation

Explanations and Expectations



OpenMP implementation

- linear speedup due to few communication
- suitable for smaller problem sizes
 - but: already used CCR's "biggest" machine

MPI implementation

- slower communication, setup overhead

Explanations and Expectations



OpenMP implementation

- linear speedup due to few communication
 - suitable for smaller problem sizes
- but: already used CCR's "biggest" machine

MPI implementation

- slower communication, setup overhead
- more processors available

Explanations and Expectations



OpenMP implementation

- linear speedup due to few communication
 - suitable for smaller problem sizes
- but: already used CCR's "biggest" machine

MPI implementation

- slower communication, setup overhead
 - more processors available
- suitable for larger problem sizes

Explanations and Expectations



OpenMP implementation

- linear speedup due to few communication
 - suitable for smaller problem sizes
- but: already used CCR's "biggest" machine

MPI implementation

- slower communication, setup overhead
 - more processors available
- suitable for larger problem sizes

compare

Explanations and Expectations



OpenMP implementation

- linear speedup due to few communication
 - suitable for smaller problem sizes
- but: already used CCR's "biggest" machine

MPI implementation

- slower communication, setup overhead
 - more processors available
- suitable for larger problem sizes

compare $x \in \{0, 1\}^{32}$

Explanations and Expectations



OpenMP implementation

- linear speedup due to few communication
 - suitable for smaller problem sizes
- but: already used CCR's "biggest" machine

MPI implementation

- slower communication, setup overhead
 - more processors available
- suitable for larger problem sizes

compare $x \in \{0, 1\}^{32}$

OpenMP

MPI

Explanations and Expectations



OpenMP implementation

- linear speedup due to few communication
 - suitable for smaller problem sizes
- but: already used CCR's "biggest" machine

MPI implementation

- slower communication, setup overhead
 - more processors available
- suitable for larger problem sizes

compare $x \in \{0, 1\}^{32}$

PEs 32

OpenMP

MPI

Explanations and Expectations



OpenMP implementation

- linear speedup due to few communication
 - suitable for smaller problem sizes
- but: already used CCR's "biggest" machine

MPI implementation

- slower communication, setup overhead
 - more processors available
- suitable for larger problem sizes

compare $x \in \{0, 1\}^{32}$

# PEs	32
-------	----

OpenMP	152s
--------	------

MPI	174s
-----	------

Explanations and Expectations



OpenMP implementation

- linear speedup due to few communication
- suitable for smaller problem sizes
but: already used CCR's "biggest" machine

MPI implementation

- slower communication, setup overhead
- more processors available
suitable for larger problem sizes

compare $x \in \{0, 1\}^{32}$

# PEs	32	64
-------	----	----

OpenMP	152s	
--------	------	--

MPI	174s	
-----	------	--

Explanations and Expectations



OpenMP implementation

- linear speedup due to few communication
- suitable for smaller problem sizes
but: already used CCR's "biggest" machine

MPI implementation

- slower communication, setup overhead
- more processors available
suitable for larger problem sizes

compare $x \in \{0, 1\}^{32}$

# PEs	32	64
OpenMP	152s	—
MPI	174s	97s

Explanations and Expectations



OpenMP implementation

- linear speedup due to few communication
 - suitable for smaller problem sizes
- but: already used CCR's "biggest" machine

MPI implementation

- slower communication, setup overhead
 - more processors available
- suitable for larger problem sizes

CUDA implementation

Explanations and Expectations



OpenMP implementation

- linear speedup due to few communication
 - suitable for smaller problem sizes
- but: already used CCR's "biggest" machine

MPI implementation

- slower communication, setup overhead
 - more processors available
- suitable for larger problem sizes

CUDA implementation (future work)

Explanations and Expectations



OpenMP implementation

- linear speedup due to few communication
 - suitable for smaller problem sizes
- but: already used CCR's "biggest" machine

MPI implementation

- slower communication, setup overhead
 - more processors available
- suitable for larger problem sizes

CUDA implementation (future work)

- fast communication between processing elements

Explanations and Expectations



OpenMP implementation

- linear speedup due to few communication
 - suitable for smaller problem sizes
- but: already used CCR's "biggest" machine

MPI implementation

- slower communication, setup overhead
 - more processors available
- suitable for larger problem sizes

CUDA implementation (future work)

- fast communication between processing elements
- very high number of processors on single nodes

Sources and References

Hans Delf and Helmut Knebl. *Introduction to Cryptography: Principles and Applications*. Springer, 2007.

Alfred J. Menezes, Paul C. van Oorschot, Scott A. Vanstone. *Handbook of Applied Cryptography*. CRC Press, 1996.

Charles P. Pfleeger and Shari L. Pfleeger. *Analyzing Computer Security: A Threat/Vulnerability/Countermeasure Approach* Prentice Hall, 2011.