Fewer Round Trips: Tail Latency in Small Replicated Stores

Mira Okafor, Tomas Reyes

Sample text for the mksite demo. The study, the data and the results in this paper are invented, so that you can see how the reader, the citation box and the Google Scholar tags behave.

1 Introduction

1 Introduction depends on how the work is measured, and on who does the measuring. In this section we describe the setting, the people and the equipment, and we explain why we chose the methods we did instead of the more common alternatives.

Every figure was checked twice by two different people, and any disagreement was settled by repeating the measurement on another day. Where a number could not be repeated, it is marked as such in the tables and left out of the summary statistics.

We found that small differences in set-up produced larger differences in the result than the quantity we were trying to study. That is not a reason to give up. It is a reason to report the set-up in more detail than is usual, and to publish the raw data so that others can look for what we missed.

The most surprising result came from the simplest comparison. When we put the two approaches side by side on the same day, in the same place, the gap between them was about half the size we expected from earlier work. We discuss possible reasons below, and we are careful not to claim more than the data support.

1 Introduction depends on how the work is measured, and on who does the measuring. In this section we describe the setting, the people and the equipment, and we explain why we chose the methods we did instead of the more common alternatives.

Every figure was checked twice by two different people, and any disagreement was settled by repeating the measurement on another day. Where a number could not be repeated, it is marked as such in the tables and left out of the summary statistics.

We found that small differences in set-up produced larger differences in the result than the quantity we were trying to study. That is not a reason to give up. It is a reason to report the set-up in more detail than is usual, and to publish the raw data so that others can look for what we missed.

2 Design

2 Design depends on how the work is measured, and on who does the measuring. In this section we describe the setting, the people and the equipment, and we explain why we chose the methods we did instead of the more common alternatives.

Every figure was checked twice by two different people, and any disagreement was settled by repeating the measurement on another day. Where a number could not be repeated, it is marked as such in the tables and left out of the summary statistics.

We found that small differences in set-up produced larger differences in the result than the quantity we were trying to study. That is not a reason to give up. It is a reason to report the set-up in more detail than is usual, and to publish the raw data so that others can look for what we missed.

The most surprising result came from the simplest comparison. When we put the two approaches side by side on the same day, in the same place, the gap between them was about half the size we expected from earlier work. We discuss possible reasons below, and we are careful not to claim more than the data support.

2 Design depends on how the work is measured, and on who does the measuring. In this section we describe the setting, the people and the equipment, and we explain why we chose the methods we did instead of the more common alternatives.

Every figure was checked twice by two different people, and any disagreement was settled by repeating the measurement on another day. Where a number could not be repeated, it is marked as such in the tables and left out of the summary statistics.

We found that small differences in set-up produced larger differences in the result than the quantity we were trying to study. That is not a reason to give up. It is a reason to report the set-up in more detail than is usual, and to publish the raw data so that others can look for what we missed.

The most surprising result came from the simplest comparison. When we put the two approaches side by side on the same day, in the same place, the gap between them was about half the size we expected from earlier work. We discuss possible reasons below, and we are careful not to claim more than the data support.

2 Design depends on how the work is measured, and on who does the measuring. In this section we describe the setting, the people and the equipment, and we explain why we chose the methods we did instead of the more common alternatives.

3 Evaluation

3 Evaluation depends on how the work is measured, and on who does the measuring. In this section we describe the setting, the people and the equipment, and we explain why we chose the methods we did instead of the more common alternatives.

Every figure was checked twice by two different people, and any disagreement was settled by repeating the measurement on another day. Where a number could not be repeated, it is marked as such in the tables and left out of the summary statistics.

We found that small differences in set-up produced larger differences in the result than the quantity we were trying to study. That is not a reason to give up. It is a reason to report the set-up in more detail than is usual, and to publish the raw data so that others can look for what we missed.

The most surprising result came from the simplest comparison. When we put the two approaches side by side on the same day, in the same place, the gap between them was about half the size we expected from earlier work. We discuss possible reasons below, and we are careful not to claim more than the data support.

3 Evaluation depends on how the work is measured, and on who does the measuring. In this section we describe the setting, the people and the equipment, and we explain why we chose the methods we did instead of the more common alternatives.

Every figure was checked twice by two different people, and any disagreement was settled by repeating the measurement on another day. Where a number could not be repeated, it is marked as such in the tables and left out of the summary statistics.

We found that small differences in set-up produced larger differences in the result than the quantity we were trying to study. That is not a reason to give up. It is a reason to report the set-up in more detail than is usual, and to publish the raw data so that others can look for what we missed.

The most surprising result came from the simplest comparison. When we put the two approaches side by side on the same day, in the same place, the gap between them was about half the size we expected from earlier work. We discuss possible reasons below, and we are careful not to claim more than the data support.

3 Evaluation depends on how the work is measured, and on who does the measuring. In this section we describe the setting, the people and the equipment, and we explain why we chose the methods we did instead of the more common alternatives.

4 Related work

4 Related work depends on how the work is measured, and on who does the measuring. In this section we describe the setting, the people and the equipment, and we explain why we chose the methods we did instead of the more common alternatives.

Every figure was checked twice by two different people, and any disagreement was settled by repeating the measurement on another day. Where a number could not be repeated, it is marked as such in the tables and left out of the summary statistics.

We found that small differences in set-up produced larger differences in the result than the quantity we were trying to study. That is not a reason to give up. It is a reason to report the set-up in more detail than is usual, and to publish the raw data so that others can look for what we missed.

The most surprising result came from the simplest comparison. When we put the two approaches side by side on the same day, in the same place, the gap between them was about half the size we expected from earlier work. We discuss possible reasons below, and we are careful not to claim more than the data support.

4 Related work depends on how the work is measured, and on who does the measuring. In this section we describe the setting, the people and the equipment, and we explain why we chose the methods we did instead of the more common alternatives.

5 Conclusion

5 Conclusion depends on how the work is measured, and on who does the measuring. In this section we describe the setting, the people and the equipment, and we explain why we chose the methods we did instead of the more common alternatives.

Every figure was checked twice by two different people, and any disagreement was settled by repeating the measurement on another day. Where a number could not be repeated, it is marked as such in the tables and left out of the summary statistics.

We found that small differences in set-up produced larger differences in the result than the quantity we were trying to study. That is not a reason to give up. It is a reason to report the set-up in more detail than is usual, and to publish the raw data so that others can look for what we missed.

The most surprising result came from the simplest comparison. When we put the two approaches side by side on the same day, in the same place, the gap between them was about half the size we expected from earlier work. We discuss possible reasons below, and we are careful not to claim more than the data support.

5 Conclusion depends on how the work is measured, and on who does the measuring. In this section we describe the setting, the people and the equipment, and we explain why we chose the methods we did instead of the more common alternatives.