




Olhão 2022

Performance Improvements in Set Operations

Karta S. Kooner

October 10, 2022






Olhãõ 2022


Performance is an important aspect of the interpreter

Many performance enhancements introduced over many years


Dyalog 18.0 introduced significant changes to the performance profile of the interpreter




1-byte int




2-byte int




4-byte int




8-byte double




1-byte int




2-byte int



4-byte int



8-byte double




DYALOG

Olhão 2022

Key questions to ask:

- what does the performance profile look like?
- are the assumptions of the past still valid?
- has the input space been sufficiently well tested?




DYALOG

Olhão 2022

Dyalog 18.2 reverted some performance enhancements introduced in Dyalog 18.0 for safety reasons

How should these be reinstated?


Should they be reinstated?




DYALOG

Olhão 2022


Good opportunity to step back and evaluate




$X \leftarrow ?1000p10*6 \diamond Y \leftarrow ?1000p10*6 \diamond X \in Y$




$X \leftarrow ?(2 \times 10^6)p3442210 \diamond Y \leftarrow ?5000p4000 \diamond X \in Y$





$X \leftarrow ?(2 \times 10^6)p3442210 \diamond Y \leftarrow ?5000p4000 \diamond X \in Y$




$X \leftarrow ?(2 \times 10^6)p3442210 \diamond Y \leftarrow ?5000p4000 \diamond X \in Y$















$X \in Y$ — Total elapsed time / ms





$X \in Y$ — Total elapsed time / ms





$X \in Y$ — Total elapsed time / ms





$X \in Y$ — Total elapsed time / ms





$X \in Y$ — Total elapsed time / ms





$X \in Y$ — Total elapsed time / ms





$X \in Y$ — Total elapsed time / ms





$X \in Y$ — Total elapsed time / ms





$X \in Y$ — Total elapsed time / ms



$X \in Y$ — Total elapsed time / ms



$X \in Y$ — Total elapsed time / ms




DYALOG

Olhão 2022

Determining the true performance profile of a primitive
is a difficult task

Sampling the input space can still be insightful




DYALOG

Olhão 2022

Reintroduce 18.0 performance improvements provided they are:

- significant (ideally, an order of magnitude₂)
- likely to be useful to the user
- adequately tested



DYALOG

Olhão 2022

Provide a mechanism to gather usage statistics

Possible target data includes:

- which algorithms are being hit, or would have been hit
- argument properties: set sizes, set range...
- timing statistics

