



Distributed Detection of Sequential Anomalies in Univariate Time Series


1. RAW TIMESERIES



2. DISTRIBUTE TIMESERIES WITH OVERLAP




3. PCA AND ROTATION ON EVERY MACHINE




4. REDUCE TO MAIN NODE AND SEGMENT




5. NODE EXTRACTION ON EVERY MACHINE



6. GRAPH AND SCORING ON EVERY MACHINE



7. REDUCE SCORES TO MAIN NODE



EVALUATION

GB	Runtime (s)		DADS 1P-1T	1P-20T	12P-20T
	S2G	S2G ⁺			
0.001	4	4	6	5	7
0.011	35	29	12	7	7
0.108	388	297	145	19	9
1.080	5 897	3 077	904	145	26
5.400	‡	15 947	4 575	727	106
10.800	†	†	†	†	206
21.600	†	†	†	†	401
32.400	†	†	†	†	586
43.200	†	†	†	†	802
54.000	†	†	†	†	986

Execution times (best of 3 runs) of DADS, S2G, and S2G⁺ when analyzing a snippet/concatenation of the MBA time series of patient 14046. † marks *memory limit of 58 GiB exceeded* and ‡ marks *execution time limit of 8 hours exceeded*.

Phillip Wenig
Johannes Schneider
Thorsten Papenbrock

Information Systems Group
Hasso Plattner Institute, Potsdam, Germany

E-Mail: firstname.lastname@hpi.de



<https://hpi.de/naumann/s/dads>