



Trace Compass, Data mining and Execution Comparison

Fateme Faraji Daneshgar

Polytechnique Montréal

DORSAL Laboratory

Agenda

- An Execution Comparison Tool
- Inter process metrics for execution comparison based on frequent pattern mining

Agenda

- An Execution Comparison Tool
 - Input Trace
 - Filtering
 - Merging Execution Call Graphs
 - Differential Flame Graph
 - Future Work
- Inter process metrics for execution comparison based on frequent pattern mining
 - Data Mining
 - Data Mining and Execution Comparison
 - Execution Clustering
 - Analysis
 - Future Work

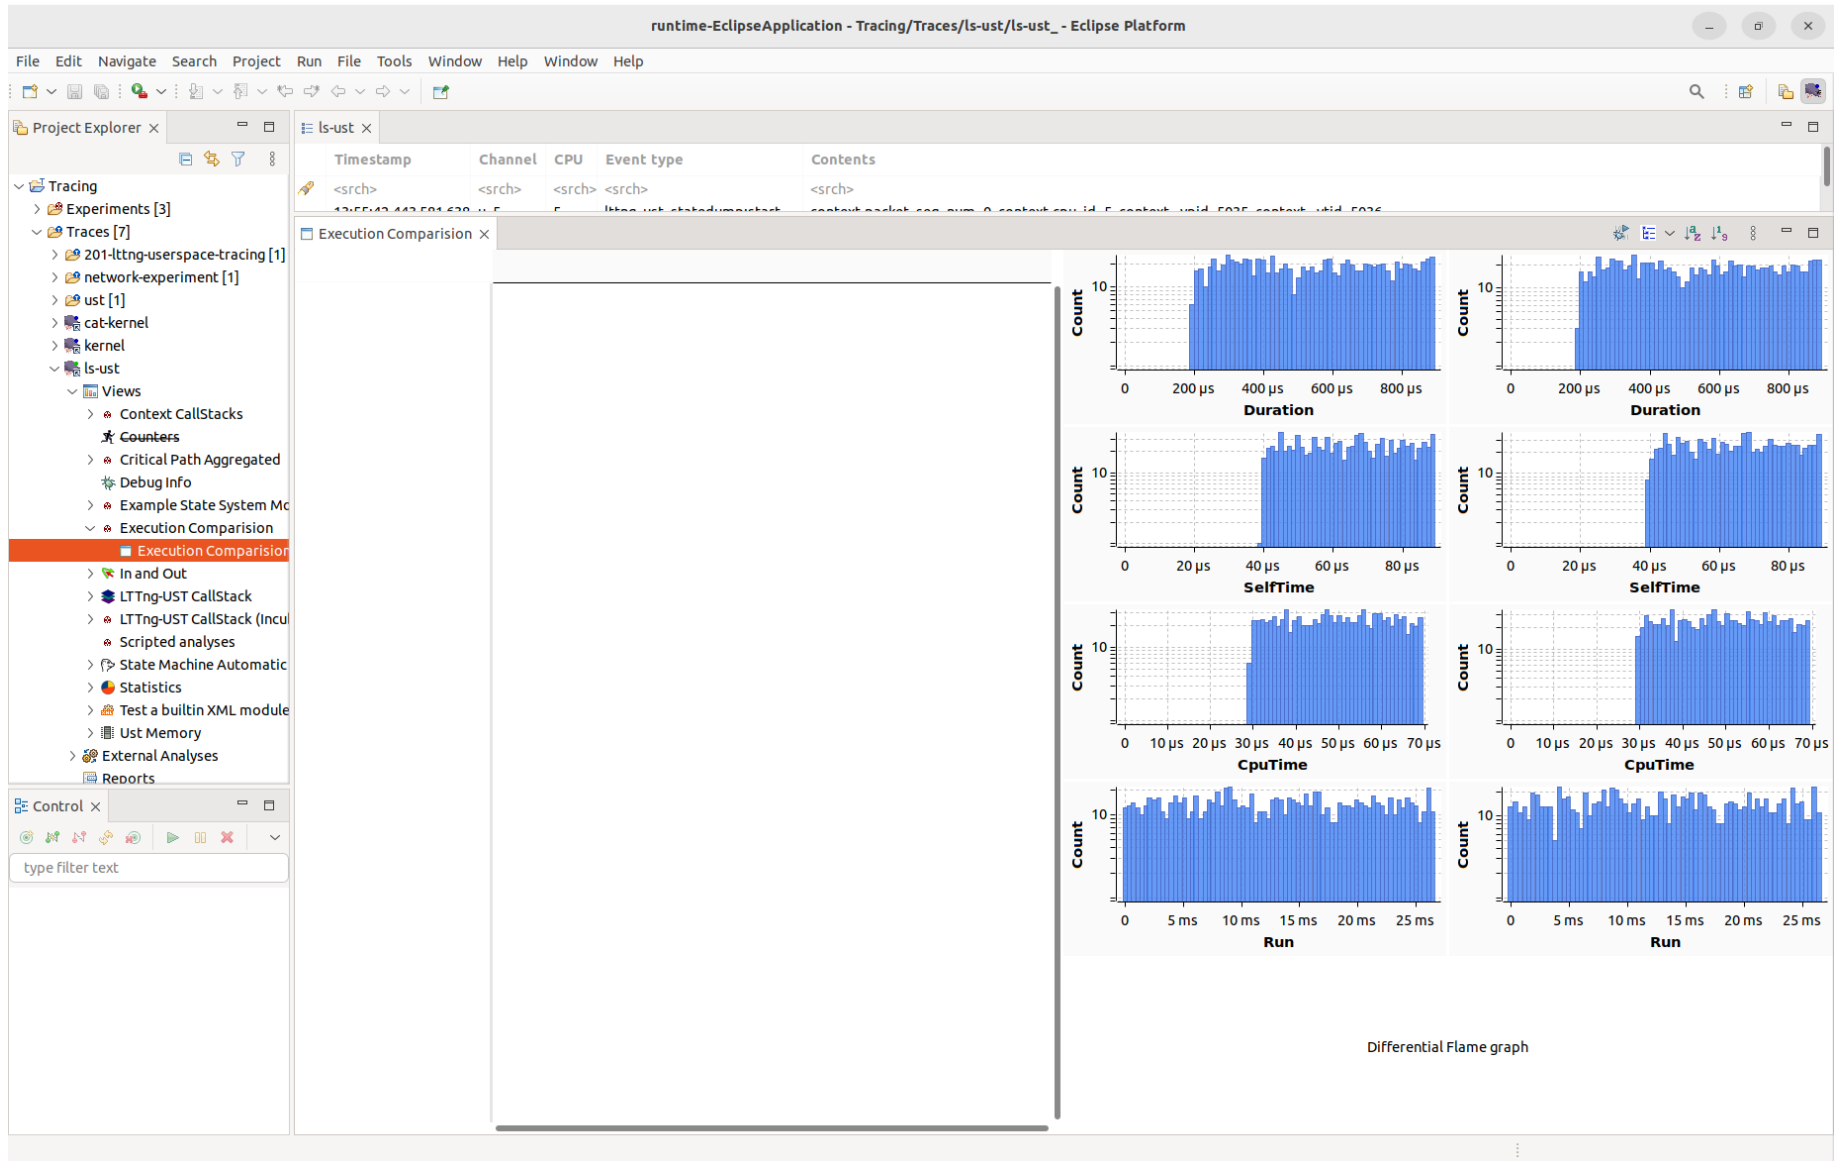
An execution comparison tool

- Use Case
 - Highlighting differences between two sets of executions of the same task
 - Fast vs slow queries
 - Fast vs slow release version
- Inspiration
 - Doray, Francois, and Michel Dagenais. "Diagnosing performance variations by comparing multi-level execution traces." IEEE Transactions on Parallel and Distributed Systems 28.2 (2016): 462-474

Input Trace

- Lttng ust trace
 - Cyg profile events
 - Each execution contained one thread
- A call graph is built for each execution
 - Density charts
 - Aggregated called function statistics
 - Duration
 - Cpu time
 - Self time
 - Run

Input Trace

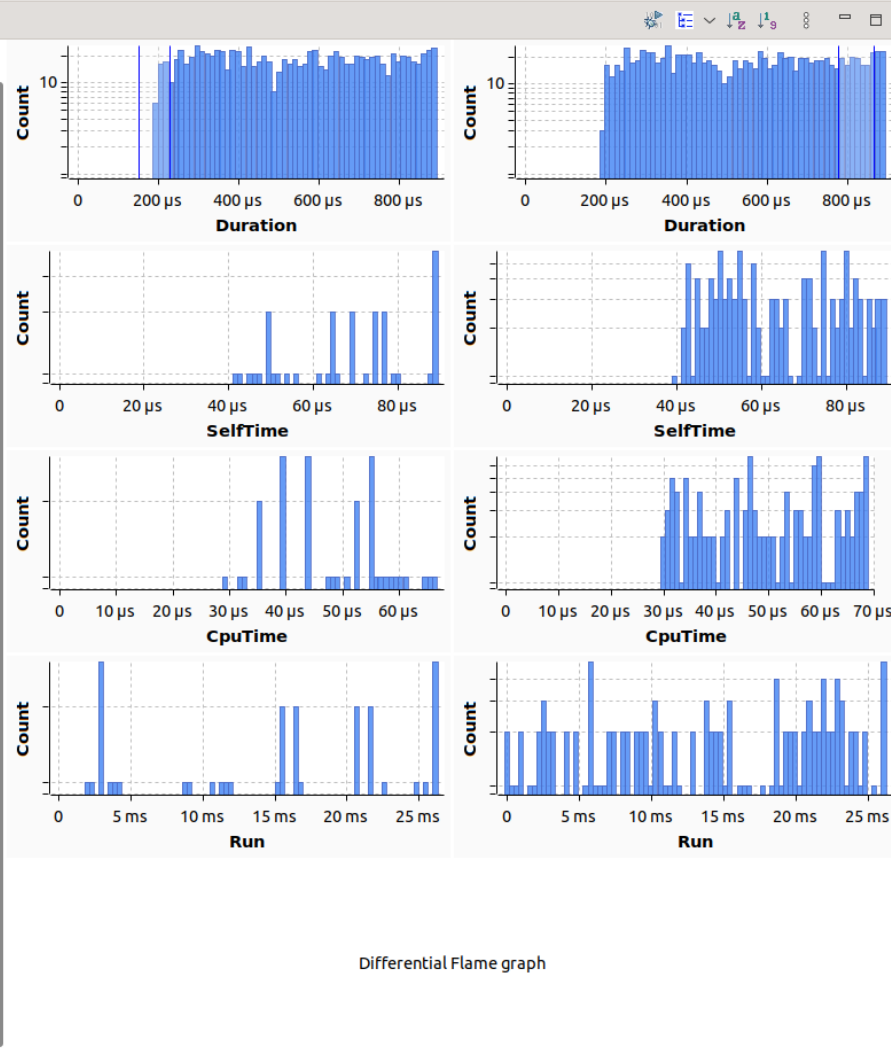



Filtering

- A group of executions can be selected in each column
- Other density charts will be updated to show the statistics of selected executions

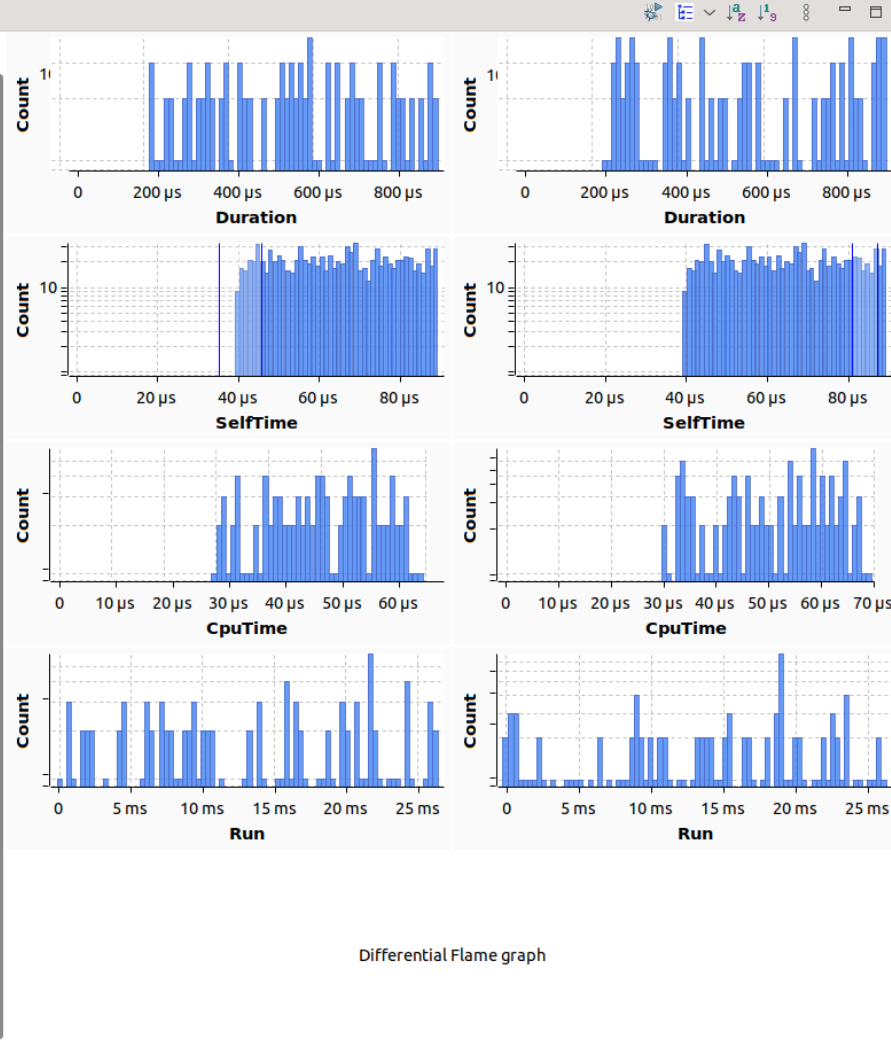

- Tracing
- Experiments [3]
- Traces [7]
 - 201-ltng-userspace-tracing [1]
 - network-experiment [1]
 - ust [1]
 - cat-kernel
 - kernel
 - ls-ust
 - Views
 - Context CallStacks
 - Counters
 - Critical Path Aggregated
 - Debug Info
 - Example State System Mc
 - Execution Comparision
 - Execution Comparision
 - In and Out
 - LTTng-UST CallStack
 - LTTng-UST CallStack (Inclu
 - Scripted analyses
 - State Machine Automatic
 - Statistics
 - Test a builtin XML module
 - Ust Memory
 - External Analyses
 - Reports

Timestamp	Channel	CPU	Event type	Contents
<srch>	<srch>	<srch>	<srch>	<srch>
12:55:42.442581620	5	5	ltng_ust_stateuserstart	context:pid: 600, sum: 0, context:cpu: id: 5, context:uid: 5025, context:uid: 5026



- Tracing
- Experiments [3]
- Traces [7]
 - 201-ltng-userspace-tracing [1]
 - network-experiment [1]
 - ust [1]
 - cat-kernel
 - kernel
 - ls-ust
 - Views
 - Context CallStacks
 - Counters
 - Critical Path Aggregated
 - Debug Info
 - Example State System Mc
 - Execution Comparision
 - Execution Comparision
 - In and Out
 - LTTng-UST CallStack
 - LTTng-UST CallStack (Inclu
 - Scripted analyses
 - State Machine Automatic
 - Statistics
 - Test a builtin XML module
 - Ust Memory
 - External Analyses
 - Reports

Timestamp	Channel	CPU	Event type	Contents
<srch>	<srch>	<srch>	<srch>	<srch>
12:55:42.442581620	5	5	ltns-ust_startuserstack	context.addr=0 context.sum=0 context.uid=5 context.uid=5036

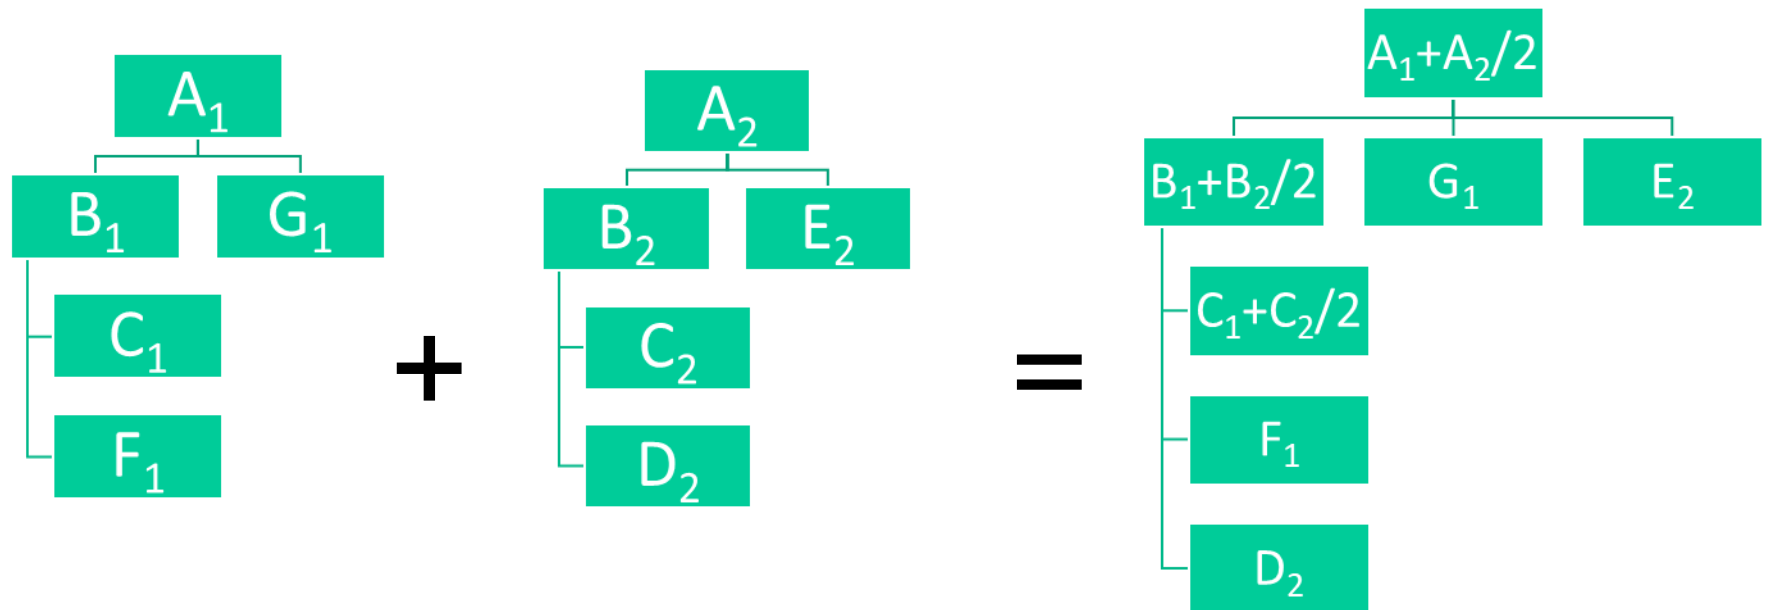


Control x

type filter text

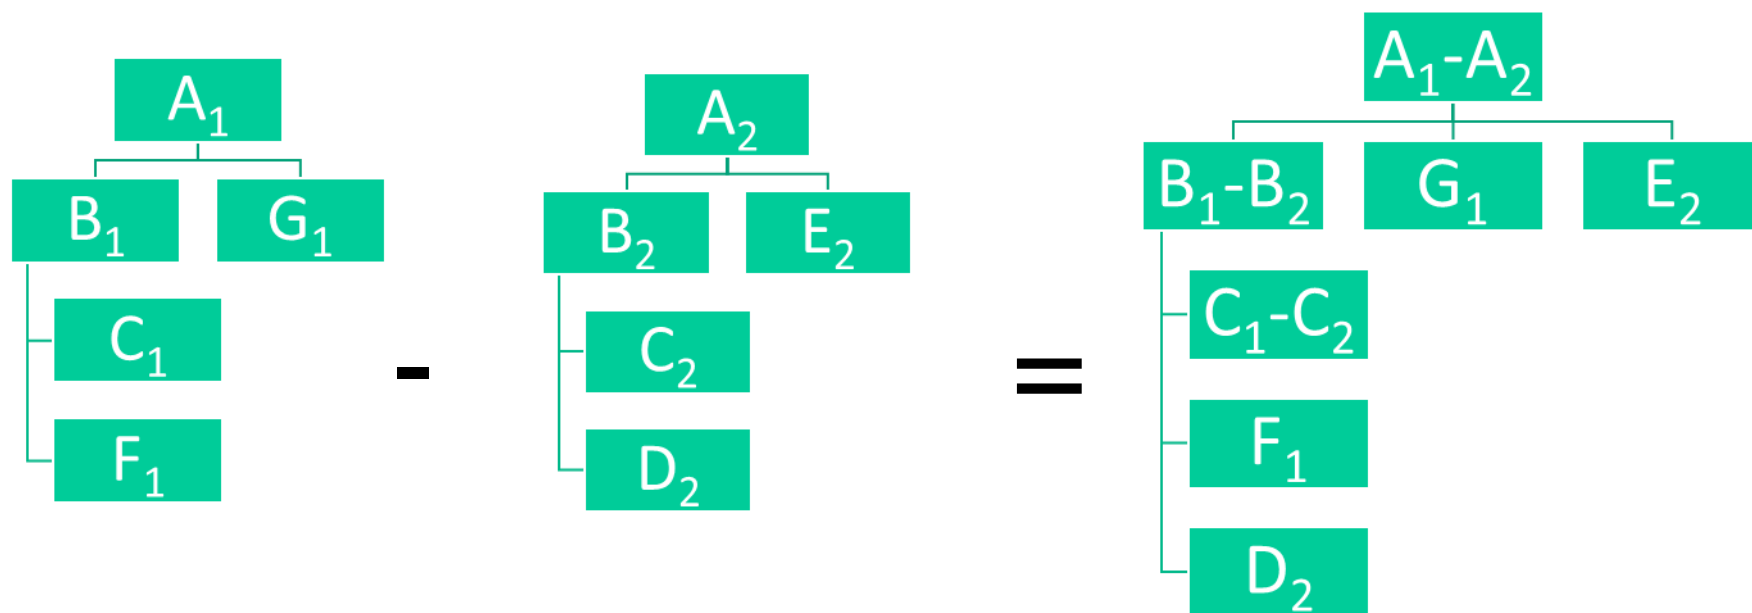
Merging Execution CallGraph

- The call graphs of the selected execution in each group are merged using Average operator



Differential Flam Graph

- The differential call graph is computed and the differential flame graph is shown



runtime-EclipseApplication - Tracing/Traces/ls-ust/ls-ust_ - Eclipse Platform

File Edit Navigate Search Project Run File Tools Window Help Window Help

Project Explorer x

- Tracing
 - Experiments [3]
 - Traces [7]
 - 201-ltng-userspace-tracing [1]
 - network-experiment [1]
 - ust [1]
 - cat-kernel
 - kernel
 - ls-ust
 - Views
 - Context CallStacks
 - Counters
 - Critical Path Aggregated
 - Debug Info
 - Example State System Mc
 - Execution Comparison
 - In and Out
 - LTng-UST CallStack
 - LTng-UST CallStack (Inclu
 - Scripted analyses
 - State Machine Automatic
 - Statistics
 - Test a builtin XML module
 - Ust Memory
 - External Analyses
 - Reports

Control x

type filter text

Timestamp	Channel	CPU	Event type	Contents
<srch>	<srch>	<srch>	<srch>	<srch>
12:55:43.443.504.620	5	5	ltng-ust-ustadumrstart	coastart-astart-esp-enum-0-coastart-uid-5-coastart-uid-5035-coastart-uid-5035

Execution Comparison x

0 10,000,000

ls-ust

5035

5035

0 0x4720 (/home/gbastien/Dorsal/sources/coreutils/src/ls)

1 0x...0x10190 (/home/gbastien/Dorsal/sources/coreutils/src/...

2 0... 0x10020 (/home/gbastien/Dorsal/sources/coreutils/...

3 0... 0xf600 (/home/gbastien/Dorsal/sources/coreutils/...

4 0x... 0x98d0 (/home/gbasti... 0x9950 (/home/gbasti...

5 0x97a0 (/home/gbasti... 0x97a0 (/home/gbasti...

6 0x1e930 (/home/gbas... 0x1e930 (/home/gbas...

7

8

9

10

11

12

Differential Flame graph

Future Work

- Integrating with Maryam's project
 - Supporting call stack for span data
 - Supporting parallel executions
 - Merge to kernel level metrics provided by Maryam
- Integrating with data mining metrics

Inter Process Execution Metrics Based on Data Mining

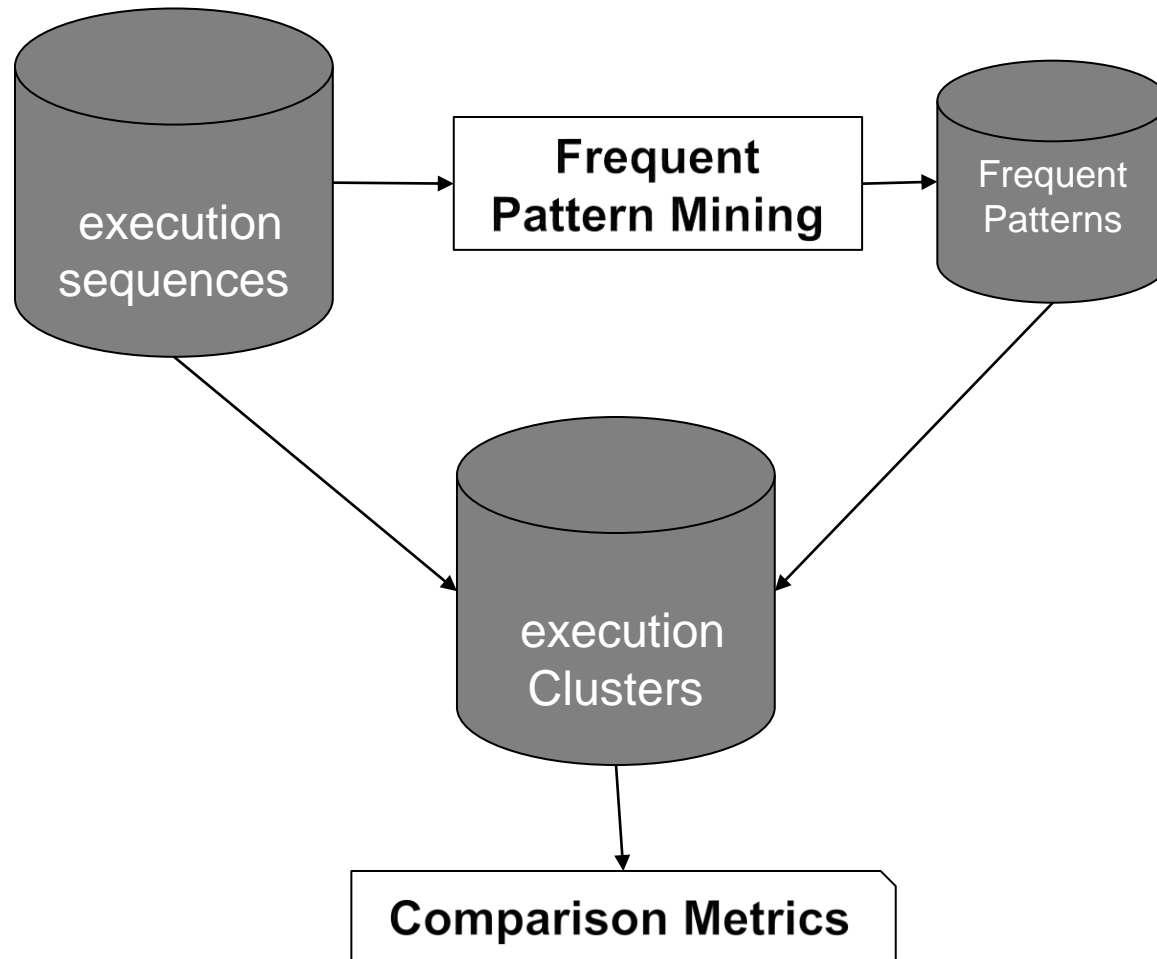
Data Mining

- Extracting knowledge from huge raw data
- To reveal the differences between executions
 - Beside process specific metrics
 - **Inter-process** knowledge hidden in software logs

Data Mining/Execution Comparision

- The similarity rate of each execution with others can be a comparison metric
- Frequent sequential pattern mining to the sequence of *system calls* in each execution
- The executions are **clustered** based on the extracted frequent patterns

Execution Clustering



Execution Clustering

P2:4

S1 = accept4,accept4,getsockname,fcntl,fcntl,read,newstat,open,open,open,open,
rt_sigprocmask,getcwd,chdir,setitimer,fcntl,access,access,....

P1:3

S1 = P2:4,P1:3,accept4,accept4,newstat,open,open,open,open,rt_sigprocmask,....

Cluster Pattern

Non-Pattern

- Each occurrence of FPs in sequences are replaced with the name of corresponding FP
- These occurrences are separated from remainder of sequence and called “Cluster Pattern” and the remainder of sequence is called “Non-Pattern”.
- Anomaly Degree: The ratio of Cluster pattern part to the Non-pattern part

Execution Metrics

- Executions are clustered based on cluster patterns
 - **NormalCnt**: the number of normal executions
 - **LongCnt**: the number of long execution
 - **DurMax**: the maximum of execution duration in a cluster
 - **DurMin**: the minimum of execution duration in a cluster
 - **DurAvg**: the average of execution duration in a cluster
 - **DurSd**: the Standard deviation of execution duration in a cluster
 - **AnoMin**: the minimum of execution anomaly degree in a cluster
 - **AnoMax**: the maximum of execution anomaly degree in a cluster
 - **AnoAvg**: the average of execution anomaly degree in a cluster
 - **AnoSd**: the Standard deviation of anomaly degree duration in a cluster

Frequent Patterns x

Pattern	NormalExec	LongExec	DurMax	DurMin	DurAvg	DurSd	AnomalyMin	AnomalyAvg	AnomalySd
NOPATTER	11	9	21809540514	143172	3.47494195E9	7.7045509E9	0.004	0.326	0.362
P6:1	13	2	21806578232	26452628	1.56895949E9	5.4105897E9	0.369	0.55	0.214
P5:2	13	2	21810851549	26226462	1.60396557E9	5.4021217E9	0.453	0.599	0.14
P1:1,P2:1,P3:2	215	3	22111733434	30066342	2.1791752E8	1.48704589E9	0.662	0.8	0.066
P5:3,P2:1,P3:2,P4:3	1	1	552305570	61026135	3.06665856E8	2.45639728E8	0.946	0.973	0.027
P3:2,P5:2,P4:2	21	3	1031558310	21562158	1.89314608E8	2.15171616E8	0.28	0.381	0.165
P3:2,P5:2	16	1	1018085956	50339593	1.8109848E8	2.13005664E8	0.561	0.663	0.114
P5:3,P2:1,P4:3	1	1	560287720	135295308	3.4779152E8	2.12496224E8	0.945	0.972	0.028
P1:1,P2:1,P4:3	5	1	591869941	66850735	1.93743968E8	1.80658336E8	0.823	0.926	0.054
P1:1,P2:1,P3:2,P4:3	6	2	541469848	87920019	2.24265664E8	1.78433328E8	0.789	0.892	0.062
P1:1,P5:2,P4:2	33	1	1093520117	26474379	1.53827296E8	1.67199872E8	0.722	0.879	0.061
P1:1,P2:1,P3:2,P5:2,P4:3	10	2	591907178	51783176	1.8361688E8	1.62291776E8	0.861	0.938	0.034
P5:2,P4:2	17	2	635441945	20496961	1.73063808E8	1.557876E8	0.37	0.472	0.142
P1:1,P3:2,P5:2,P4:2	45	1	1015444294	45353904	1.37624608E8	1.36713936E8	0.516	0.638	0.075
P5:3,P3:2	23	1	678306318	74500119	1.653964E8	1.10007456E8	0.802	0.895	0.052
P2:1,P3:2,P5:2	33	1	719197176	29520461	1.36168128E8	1.09094352E8	0.801	0.922	0.049
P5:2,P6:1	18	1	508671782	25853688	9.921876E7	1.0897648E8	0.531	0.609	0.103
P1:1,P4:2	100	1	1104604531	22189918	1.29870392E8	1.04778936E8	0.355	0.432	0.081
P1:1	77	3	626482150	22427779	1.33582656E8	1.02264912E8	0.57	0.686	0.068
P1:1,P2:1,P3:2,P6:4	1	1	318507770	138742644	2.286252E8	8.9882568E7	0.788	0.894	0.106
P3:2,P6:1	3	1	324039031	110607839	1.77788464E8	8.737984E7	0.879	0.917	0.048
P5:3,P2:1	76	2	617597000	31648032	1.36434544E8	8.4132184E7	0.72	0.837	0.059
P1:1,P2:1,P3:1	26	1	496109195	36227681	1.24487496E8	8.1862616E7	0.738	0.888	0.066
P5:3	107	1	315170519	22691016	8.2592912E7	5.9257272E7	0.557	0.628	0.059
P5:3,P6:1	272	0	198170596	20814824	7.091584E7	5.5133704E7	0.699	0.779	0.066
P5:3,P4:3	4	0	179302380	32669827	1.2220932E8	5.5102396E7	0.801	0.909	0.072
P4:2	5	0	196118815	34690501	1.40699104E8	5.5097776E7	0.567	0.712	0.152
P1:1,P2:1,P4:1	2	0	196295811	91920303	1.44108064E8	5.2187752E7	0.991	0.995	0.005
P1:1,P2:1,P3:2,P5:2	126	1	486502184	30629441	1.06379728E8	5.0613352E7	0.689	0.869	0.078
P1:1,P5:2,P6:1	24	1	338245593	75567994	1.42259536E8	4.8485716E7	0.748	0.864	0.068
P1:1,P5:2	23	0	192146036	22678965	1.29956144E8	4.8273148E7	0.716	0.857	0.072
P5:3,P3:2,P4:2	55	1	353415497	34181897	1.3911448E8	4.7784284E7	0.722	0.854	0.068
P2:1,P5:2	10	0	198895289	50515718	1.40456256E8	4.7765948E7	0.77	0.883	0.062
P5:3,P2:1,P3:1,P4:2	19	1	300790783	76853730	1.32948776E8	4.7757836E7	0.704	0.844	0.074
P2:1,P3:2,P5:2,P6:4	2	0	178584145	83495652	1.31039896E8	4.7544248E7	0.917	0.958	0.042
P1:1,P3:2,P5:2,P6:1	31	0	192302176	22437077	1.09232136E8	4.7343176E7	0.717	0.879	0.065
P2:1,P3:2	4	0	198527562	66209059	1.35571696E8	4.6891568E7	0.884	0.949	0.052

Analysis

- NoPattern Cluster
 - The executions that are totally different from most of the executions

Pattern	NormalExec	LongExec	DurMax	DurMin	DurAvg	DurSd	AnomalyMin	AnomalyAvg	AnomalySd
P6:1	13	2	21806578232	26452628	1.56895949E9	5.4105897E9	0.369	0.55	0.214
P5:2	13	2	21810851549	26226462	1.60396557E9	5.4021217E9	0.453	0.599	0.14
P1:1,P2:1,P3:2	215	3	22111733434	30066342	2.1791752E8	1.48704589E9	0.662	0.8	0.066
P5:3,P2:1,P3:2,P4:3	1	1	552305570	61026135	3.06665856E8	2.45639728E8	0.946	0.973	0.027
P3:2,P5:2,P4:2	21	3	1031558310	21562158	1.89314608E8	2.15171616E8	0.28	0.381	0.165
P3:2,P5:2	16	1	1018085956	50339593	1.8109848E8	2.13005664E8	0.561	0.663	0.114
P5:3,P2:1,P4:3	1	1	560287720	135295308	3.4779152E8	2.12496224E8	0.945	0.972	0.028

Analysis

- High DurSd and high AnoSd indicates the execution that although they have some similarity but their differences are considerable.

Latency cause: there are more system calls in their executions

Pattern	NormalExec	LongExec	DurMax	DurMin	DurAvg	DurSd	AnomalyMin	AnomalyAvg	AnomalySd
NOPATTER	11	9	21809540514	143172	3.47494195E9	7.7045509E9	0.004	0.326	0.362
P5:2	13	2	21810851549	26226462	1.60396557E9	5.4021217E9	0.453	0.599	0.14
P1:1,P2:1,P3:2	215	3	22111733434	30066342	2.1791752E8	1.48704589E9	0.662	0.8	0.066
P5:3,P2:1,P3:2,P4:3	1	1	552305570	61026135	3.06665856E8	2.45639728E8	0.946	0.973	0.027
P3:2,P5:2,P4:2	21	3	1031558310	21562158	1.89314608E8	2.15171616E8	0.28	0.381	0.165
P3:2,P5:2	16	1	1018085956	50339593	1.8109848E8	2.13005664E8	0.561	0.663	0.114
P5:3,P2:1,P4:3	1	1	560287720	135295308	3.4779152E8	2.12496224E8	0.945	0.972	0.028

Analysis

- High DurSd and low AnoSd indicates the execution that are similar in terms of their system call sequence.
 - Latency cause might be related to the longer duration of some system calls
- Low DurSd
 - Normal/long executions that are the same in terms of duration
 - The cluster patterns indicate their behavioral pattern

Future Work

- Adding analysis of the duration of each system call in clusters with high DurSd and low AnoSd