DAMNIT architecture

Thomas Kluyver, European XFEL 26th June 2025

🤎 Table View

Run table Context file

Plot Time comment Variables: 2 tags Filters (0) *

	Timestamp	mono_1_energy	undu_energy	BAM delay	PPODL	Pump pattern	Scan type	10	lf	10 peaks: fig	If peaks: fig	Absorption	XAS pp	XAS scan
238	19:02:28 16/06/2025	6.5400	6.4729	-143.1926	-157	train	mono							
239	19:13:18 16/06/2025	6.5400	6.4729	-149.6219	-157	train	mono						alm 1	
240	19:24:08 16/06/2025	6.5400	6.4729	-157.5511	-157	train	mono						and the	
241	19:56:09 16/06/2025	6.5550	6.4882	-176.6944	-157	pulse	mono						-ta-	
242	20:07:32 16/06/2025	6.5590	6.4922	-183.4412	-157	pulse	INCOUPLIN							
243	20:15:58 16/06/2025	6.5590	6.4922	-187.6522	-157	pulse	INCOUPLIN							
244	20:28:08 16/06/2025	6.5590	6.4922	-190.0079	-157	pulse	INCOUPLIN							
245	20:32:34 16/06/2025	6.5590	6.4922	-190.5333	-157	pulse	INCOUPLIN							
246	20:38:06 16/06/2025	6.5550	6.4881	-195.1545	-157	pulse	mono							Meg.
247	20:48:32 16/06/2025	6.5550	6.4881	-196.2476	-157	pulse	mono		H. S. S. S.				1	
248	21:01:25 16/06/2025	6.5600	6.4936	-199.1470	-162.2367	pulse	ppl odl							
249 If peaks: fig (run 257)										(aug. 257)				
250	0 XAS scan (run 257)											~ _ U ~		
251	If digitizer, 30 peaks First peak. WL:-10, WR:20													
252	300										0.004			
253	53 250 250 250 250 250 250 250 250 250 250								0.6 - unpumped	/ m/V	7	- 0.002		
254	- 200			_	200				N. A.P.		0.5 - W	AI	tr	- 0.000
255	[arb r			[arb u.							0.4 -		1	0.002 la
256	2) 150 test			tensity							0.3 -		V	0.004 9 5
257	⊆ 100 -			<u> </u>	100						0.2 -			0.006 Juli
258	50				50							/ If		0.008
259	0				0			-			0.1	V		-0.010
Doul	0.0	0.2 0.4	0.6 0	1.8 1.0	12500 12600	12700 12800 12900	13000 13100 1320	0 13300			6.53 6.54	6.55 6.56 Motor positions (mon	6.57 6.58 o)	l ded)
	L		samples	160		Samples			$11 \sim$					ica)
										(r	Display hover :	annotations

^ _ 🗆 X

9

 $\blacksquare \ \leftarrow \ \rightarrow \ \mathbf{C}$

DAMNIT! probing ultratast electron-polaron dynamics in lithium-ion battery cathodes

o

Web interface

Table

	Run	undu_energy	BAM delay	PPODL	Pump pattern	Scan type	10	lf	Laser intensity	I0 peaks: fig	If peaks: fig	Absorption	XAS pp	XAS scan	JF burst mode	JF 1M preview	JF ROI pre
134	136	6.48	-69.1	-157	pulse	mono				<u>Inil C</u>				(JAN)	0		
135	137	6.48	-77.4	-157	pulse	mono							-	<u>en a</u>	0		
136	138	6.49	-88.9	-157	pulse	INCOUPLING_							ingia -		0		
137	139	6.49	-95.2	-157	pulse	INCOUPLING_									0		
138	140	6.49	-99.9	-157	pulse	INCOUPLING_									0		
139	141	6.49	-107.1	-157	pulse	INCOUPLING_									0		
140	142	6.49	-117.1	-157	pulse	INCOUPLING_						E			0		
141	143	6.49	-127.5	-157	pulse	INCOUPLING_				MALC					0		
142	144	6.49	-134.4	-157	pulse	INCOUPLING_							-		0		
143	145	6.48	-142.5	-160	pulse	mono								PME -	0		
144	146	6.48	-150.0	-157	pulse	mono							-	The second secon	0		
145	147	6.51	-158.1	-157	pulse								1	COMP.	0		
146	148	6.51	-159.1	-157	pulse								(. 104)		0		
147	149	6.51	-159.3	-157	pulse								1		0		
148	150	6.46	-166.9	-157	pulse	mono									0		
149	151	6.47	-1 B4M	: data delay -157	train	mono							1	153	0		
150	152	6.47	-173.3	-157	train	mono				M ATU			EL.	171	0		
151	153	6.47	-177.6	-157	train	mono							1		0		
152	154	6.47	-180.0	-157	train	mono							1	57	0		
153	155	6.47	-181.5	-157	train	mono							-	t.J	0		
154	156	6.47	-179.9	-157	train	mono								- 7.7 -	0		
155	157	6.47	-178.7	-157	train	mono								1.7	0		
156	158	6.47	-180.1	-157	train	mono								5	0		

```
Data And Metadata iNspection Interactive Thing
```



Valid context.

Save Validate

$\mathbf{\frown}$		
Con	lext	le

Double-click on a cell to inspect results.	Getting update
Reload from disk	

Getting updates (f40ec1ed0353a430f9d94d2c552dfaee5a401dcd)

^ _ 🗆 X

9

.

•

Running the code



Directory structure

<proposal>/usr/Shared/amore



API

```
from damnit import Damnit
```

```
db = Damnit(7925)
```

```
data = db[100, "i0"].read()
```

https://damnit.readthedocs.io/en/latest/api/

Automate initial analysis steps, retrieve data for further analysis

Services

- Current / classic model:
 - Per-proposal listener
- New model:
 - Web API server (with GPFS access)
 - Web UI server (accessible from outside)
 - Central listener for all proposals
 - Central DB writer?

Web servers

- UI server
 - Serve JS files (compiled from Typescript)
 - Proxy to API server
- API server
 - GraphQL
 - Python & FastAPI
 - Mutual TLS (mTLS) allows requests only via UI server

Possible service accounts

- xdamnear (listener)
 - Broad permissions, no user code
 - Add to all upex Slurm reservations
- xdamnrun
 - Executes user code
 - Minimal base permissions
 - Sandboxing adds proposal access as needed
- xdamnweb (web API server)
 - Access to every proposal's DAMNIT directory

Web server interactions



Web authorisation

- Sign in via keycloak
- Look up groups by username in LDAP
 - e.g. 60006736-dmgt, fxedata
- Allow access only to corresponding proposals
- Future option: check user's permissions in myMdC

Sandboxing

- Goal: restrict context file code to 1 proposal
- xdamnear (full access) starts process as
 xdamnrun (no access)
- Process adds itself to proposal group
- Sets 'no new privileges' flag
- Executes context file

Sandboxing detail

- \$ proposal-sandbox 7925 -- some-command --foo
 - Run some-command with access to proposal 7925
 - Uses setgroups() needs CAP_SETGID capability
 - seteuid() & setegid() to switch user
 - prctl() to set no_new_privs flag

https://git.xfel.eu/dataAnalysis/proposal-sandbox/-/merge_requests/1

Sandboxing & Slurm

- Context file code can (& does) submit Slurm jobs
- Jobs would have no proposal access by default
 - Fail safe but still fails
- Need to re-enter the same sandbox in Slurm job
 - Secure way to know proposal submitter belongs to?
 - MUNGE token?
 - Use Slurm task prolog?
 - Tweak sandboxing to change primary group? (passes through Slurm)

Online processing?

- Faster results in DAMNIT are desirable
- Online processing is interesting
- Obstacles
 - No Slurm on online cluster
 - Not safe to modify SQLite DB in synced usr

DB writer coordination

- Summary data in SQLite DB for each proposal
 - No need for a DB server
 - Data all together in proposal folder
- But write contention can be an issue
 - We limit concurrent jobs to mitigate this
- Centralised writer service?
 - Could batch writes when busy

HDF5 writer coordination

- Larger results stored in HDF5 file for each run
 - Concurrent writes less likely
 - Avoid unnecessary sending/copying of data
- Concurrent writes can corrupt files
 - HDF5 uses flock(), not effective across nodes
- Prevent concurrent jobs for 1 run?
- Separate output file for each job?
 - Recombine files later on when things are quiet?