

# Navigational Data Management

A general approach to representation and exploitation of relationships in scientific data sets.

Joshua Stillerman  
Martin Greenwald  
Jason Stillerman  
John Wright

MIT Plasma Science and Fusion  
Center




Alcator C-Mod Runs, Shots and Annotations  
(**one month**)

# “I have a system for storing my data and getting it back, aren't I done?”

- Collecting data has never been easier, but...
- We're struggling to keep up with the rapidly growing volume and complexity of scientific data.

## Our Thesis

- The challenge is all about giving this mountain of data meaning and putting it into context
  - Context is about metadata and relationships among data objects - “**navigational metadata**”
  - This is not specific to one science domain
- 
- In general, our approach to capturing and exploiting this class of metadata has been ad hoc and inadequate
  - This hampers data discovery and the ability to assemble coherent, complete, useful data sets.



# Discovering and Understanding Data Is Largely About Context

---

- Context is metadata about relationships between data
- Data discovery relies on “adjacency” to find other interesting data
- In the more distant past when things were smaller and simpler, we could keep that context in our heads
  - or in our colleague's heads
- Historically we’ve each build a set of ad-hoc, domain specific tools to store, explore, and retrieve this relationship metadata.
- Similar issues confront all data intensive areas of research.
- Can we solve these problems in our own domain?
- Can we generalize these to provide solutions across a broader set of domains?

## Organizing knowledge is an old problem



## Complexity: What Sorts Of Data Might Exist From A Typical Experiment?

---

- Hierarchical data stores with raw and processed data ( $\sim 10^5$  named data objects per shot)
- Relational databases with “high level” results
- Electronic logbooks & annotation
- Experimental proposals
- Run Plans & Summaries
- Data provenance systems
- Data catalogs
- Data dictionaries
- Information about experimental campaigns & plans
- Publications & presentations
- Information about researchers, authors
- Simulation inputs & outputs
- Source code management systems
- Facility information, with details of experiment, measurement systems
- Document, drawing management systems
- QA, QC information
- WBS for projects

# All Of Those Data Are Linked In Multiple and Complex Ways

Alcator C-Mod Mini-Proposal	
MP No. 831	
Subject:	Race to Midnight: 300 kJ or Bust
From:	The Alcator C-Mod Team
Group:	All of Them
Date:	September 19, 2016
Approved by:	Date Approved:
<hr/>	
<b>1. Purpose of Experiments</b>	
<small>Include immediate goal of the experiments, scientific importance and/or programmatic relevance. Refer to any relevant program milestones.</small>	
<p>This goal of this experiment is to close out operations of Alcator C-Mod by pushing operational space to break the record in stored energy and volume averaged plasmas pressure. This further demonstrates the capabilities of compact, high-field tokamaks.</p>	
<b>2. Background</b>	
<small>Discuss Physics Basis of the proposed research. Prior results at Alcator or elsewhere, and any related work being carried out separately.</small>	
<p>While C-Mod currently holds the records the record volume averaged pressure, it's clear the machine has not fully demonstrated its capabilities. Recently, the stored energy record was broken in EDA H-mode on 1160718013 (5.4 T, 1.3 MA, nI04 ~ 1.1), reaching 253 kJ at only 4 MW of input power. Additionally, MP 826 had identified a transient high, <math>H_{98} \sim 2.0</math>, operating regime that could be taken to higher current and power to increase the stored energy. Thus, we have yet to convincingly, in earnest, demonstrate</p>	

# All Of Those Data Are Linked In Multiple and Complex Ways

Alcator C-Mod			
Select Calendar Year			
2017 2016 2015 2014 2013 2012 2011 2010 2009 2008 2007 2006 2005 2004 2003 2002 2001 2000 1999 1998 1997 1996			
Run ID	Date	MP	Description
1160930	Sep 30 2016	813	EDA Access at 8 tesla
		831	Race to Midnight: 300 kJ or Bust
		831	Race to Midnight: 300 kJ or Bust
1160929	Sep 29 2016	750	Investigation of the mode structure of the WCM with a scanning Mirror Langmuir probe
		787	Improving stability of non-inductive LHCD discharges
		832	Fast time resolution LH power deposition to SOL
1160928	Sep 28 2016	627	$\rho^*$ Dependence of Intrinsic Rotation
		823	Measurement of lower hybrid wave power using two toroidally-separated probe arrays
1160927	Sep 27 2016	828	Documenting the effect of divertor geometry on upstream scrape-off layer profiles
1160926	Sep 26 2016	815	Characterization of ICRF antenna: electrical performance, impurity contamination and SOL interaction
1160923	Sep 23 2016	759	PDI and MSE measurements with high Te helium target
		762a	I-mode thresholds and operating window at 8 T (Rev 2)
		818	Electron temperature profile stiffness in L, H, I-mode plasmas
		824	Active suppression of PDI by steepening LH launcher density profile with D port ICRF
1160922	Sep 22 2016	830	Seeking a super H-mode pedestal on C-Mod
1160921	Sep 21 2016	597	ICRF Mode Conversion Flow Drive at 8 Tesla
		829	H-mode access with an open flat plate divertor
1160920	Sep 20 2016	825	Non-dimensional parameter scans in I-mode
1160919	Sep 19 2016	597	ICRF Mode Conversion Flow Drive at 8 Tesla
1160916	Sep 16 2016	759	PDI and MSE measurements with high Te helium target
1160915	Sep 15 2016	727	Localization of the WCM and the QCM in the Er well
		827	Transition dynamics and thresholds in near DN configurations
1160914	Sep 14 2016	443a	Critical temperature gradient scale length measurements in L mode discharges
		824	Active suppression of PDI by steepening LH launcher density profile with D port ICRF

MP No. 831

atic relevance. Refer to any relevant

Alcator C-Mod by pushing  
1 volume averaged plasmas  
ct, high-field tokamaks.

re, and any related work being carried

veraged pressure, it's clear  
ntly, the stored energy  
.3 MA,  $n_{104} \sim 1.1$ ), reaching  
ad identified a transient  
r current and power to  
v in earnest demonstrate

# All Of Those Data Are Linked In Multiple and Complex Ways

Select Calendar Year			
2017 2016 2015 2014 2013 2012 2011 2010 2009 2008 2007 2006 2005 2004 2003 2002 2001 2000			
Run ID	Date	MP	Description
1160930	Sep 30 2016	813	EDA Access at 8 tesla
		831	Race to Midnight: 300 kJ or Bust
		831	Race to Midnight: 300 kJ or Bust
1160929	Sep 29 2016	750	Investigation of the mode structure of the WCM with a scanning Mirror Langmuir probe
		787	Improving stability of non-inductive LHCD discharges
		832	Fast time resolution LH power deposition to SOL
1160928	Sep 28 2016	627	rho* Dependence of Intrinsic Rotation
		823	Measurement of lower hybrid wave power using two toroidally-separated probe arrays
1160927	Sep 27 2016	828	Documenting the effect of divertor geometry on upstream scrape-off layer profiles
1160926	Sep 26 2016	815	Characterization of ICRF antenna: electrical performance, impurity contamination and SOL
1160923	Sep 23 2016	759	PDI and MSE measurements with high Te helium target
		762a	I-mode thresholds and operating window at 8 T (Rev 2)
		818	Electron temperature profile stiffness in L, H, I-mode plasmas
		824	Active suppression of PDI by steepening LH launcher density profile with D port ICRF
1160922	Sep 22 2016	830	Seeking a super H-mode pedestal on C-Mod
1160921	Sep 21 2016	597	ICRF Mode Conversion Flow Drive at 8 Tesla
		829	H-mode access with an open flat plate divertor
1160920	Sep 20 2016	825	Non-dimensional parameter scans in I-mode
1160919	Sep 19 2016	597	ICRF Mode Conversion Flow Drive at 8 Tesla
1160916	Sep 16 2016	759	PDI and MSE measurements with high Te helium target
1160915	Sep 15 2016	727	Localization of the WCM and the QCM in the Er well
		827	Transition dynamics and thresholds in near DN configurations
1160914	Sep 14 2016	443a	Critical temperature gradient scale length measurements in L mode discharges
		824	Active suppression of PDI by steepening LH launcher density profile with D port ICRF

1160930041	jwhughes	SESSION_LEADER	11:29:06:267PM
Shot 41			
We are attempting to shorten up the pre-H period of RF			
Also increased the N2 puff in the beginning			
Data system hang, followed by no power shot			

1160930042	marmar	PHYSICS_OPERATOR	Sep 30 2016 11:37:13:997PM
Ip to 1.4 MA			

1160930042	jwhughes	SESSION_LEADER	Sep 30 2016 11:57:45:183PM
Shot 42			
After a long delay, got a 1.4MA shot! 235kJ			
We're going to try for one more!			

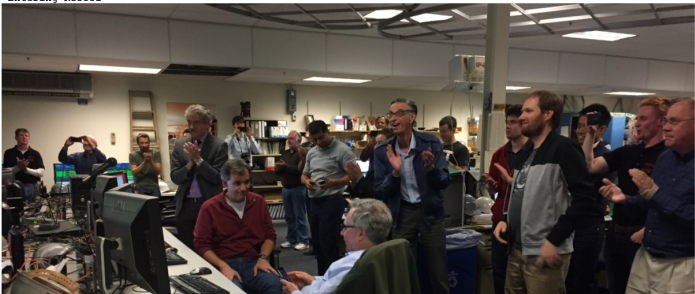
1160930043	marmar	PHYSICS_OPERATOR	Oct 1 2016 12:05:23:313AM
nl 04 to 9e19			
nitrogen to 90 msec (from 70)			
plasma, disrupt in rampdown at 1.3 seconds			
that's it folks!			

1160930043	jwhughes	SESSION_LEADER	Oct 1 2016 12:06:46:987AM
Shot 43			
That's all folks			
Matt takes the crown for stored energy.			
Best run . . . ever.			

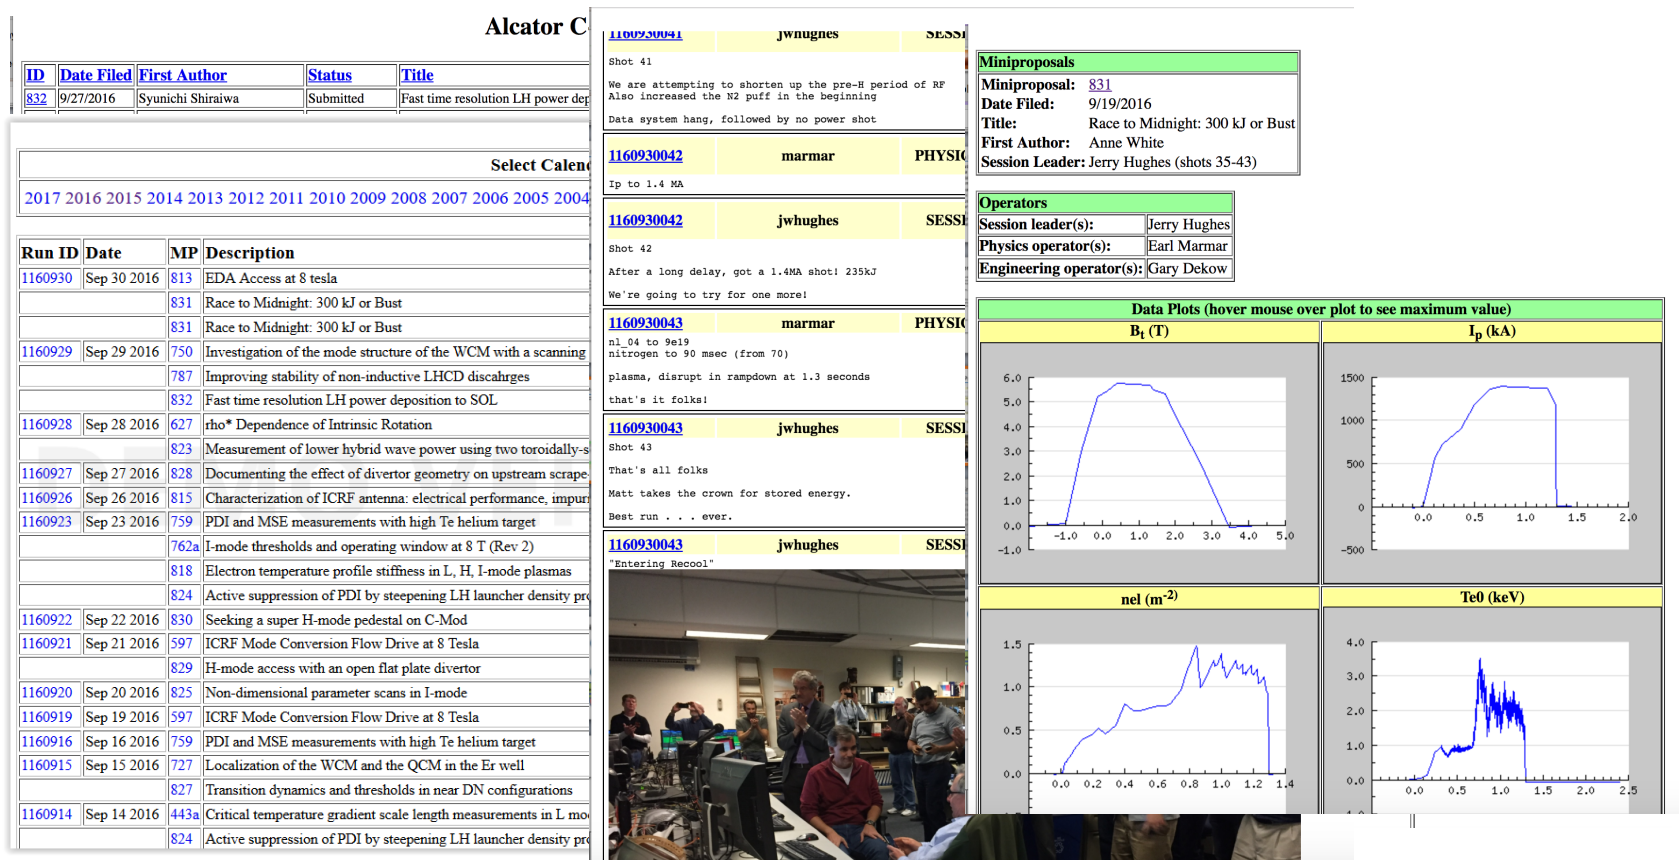
  

1160930043	jwhughes	SESSION_LEADER	Oct 5 2016 12:19:48:210PM
"Entering Recool"			

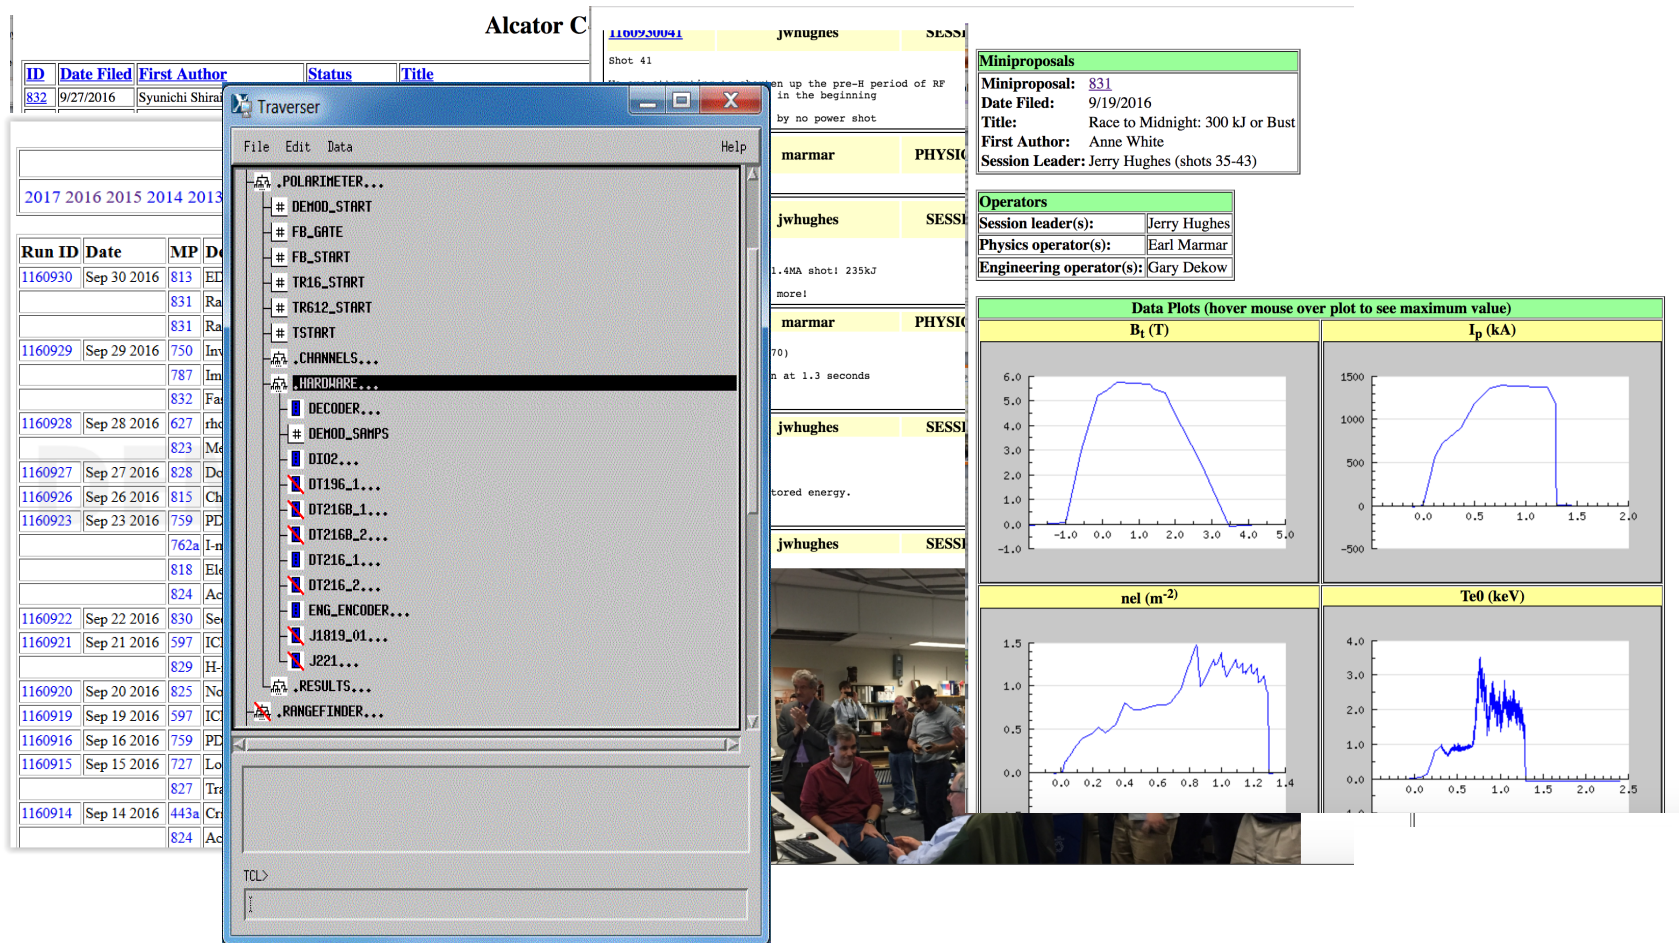


# All Of Those Data Are Linked In Multiple and Complex Ways





# All Of Those Data Are Linked In Multiple and Complex Ways

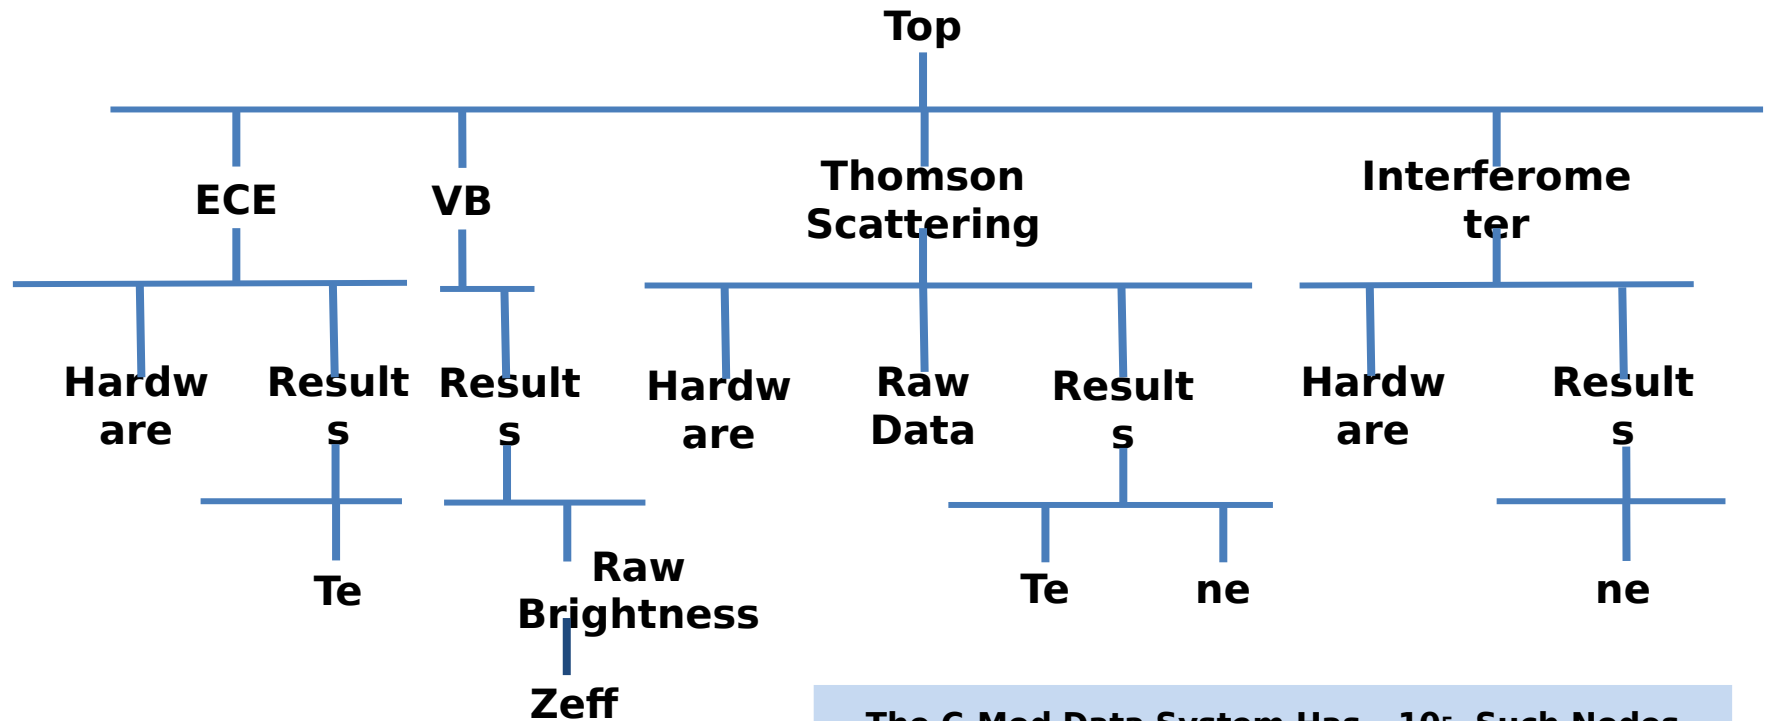


# Relationship Web Is Incomplete, Ad Hoc, Asymmetric, Singularly Organized

---

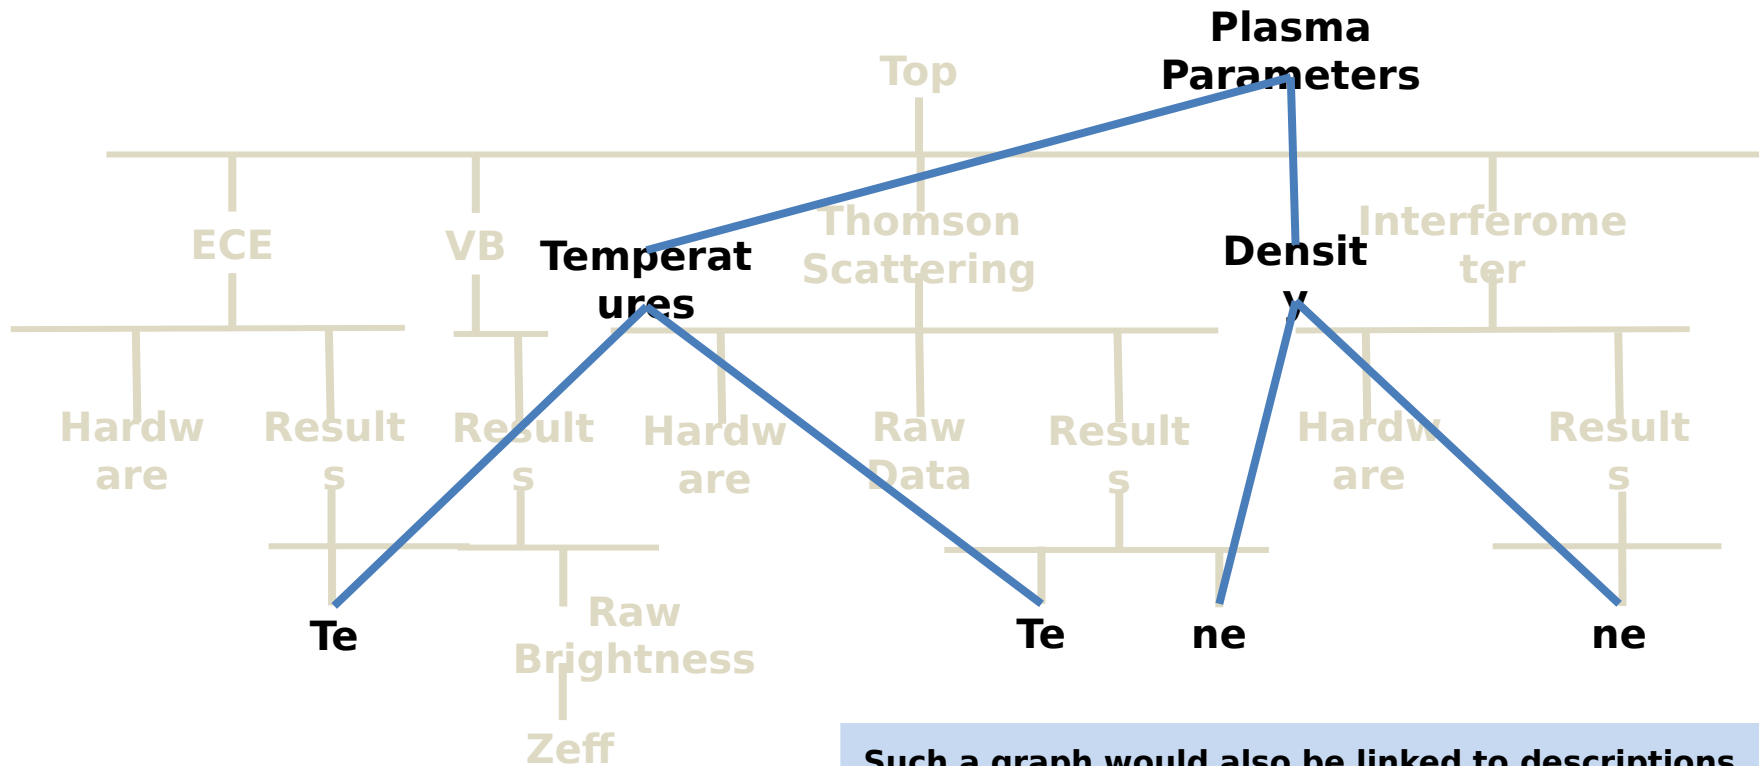
- Incomplete
  - Some relationships are explicitly represented in databases
  - Some are implicit in data or text
  - Some are only known by particular users
  - Some are not recorded and are lost forever
- Ad Hoc
  - We've added this information as needs arise
  - Schemas, vocabulary are not always consistent
  - Level of detail is uneven
- Asymmetric
  - Example: We point to interesting data from the logbook (annotation); but do not point to annotation from data (many, many other examples)
- Singularly organized
  - Trees, Tables (columns, indices), Directory/File Names

# Organization of Data - By Diagnostic System



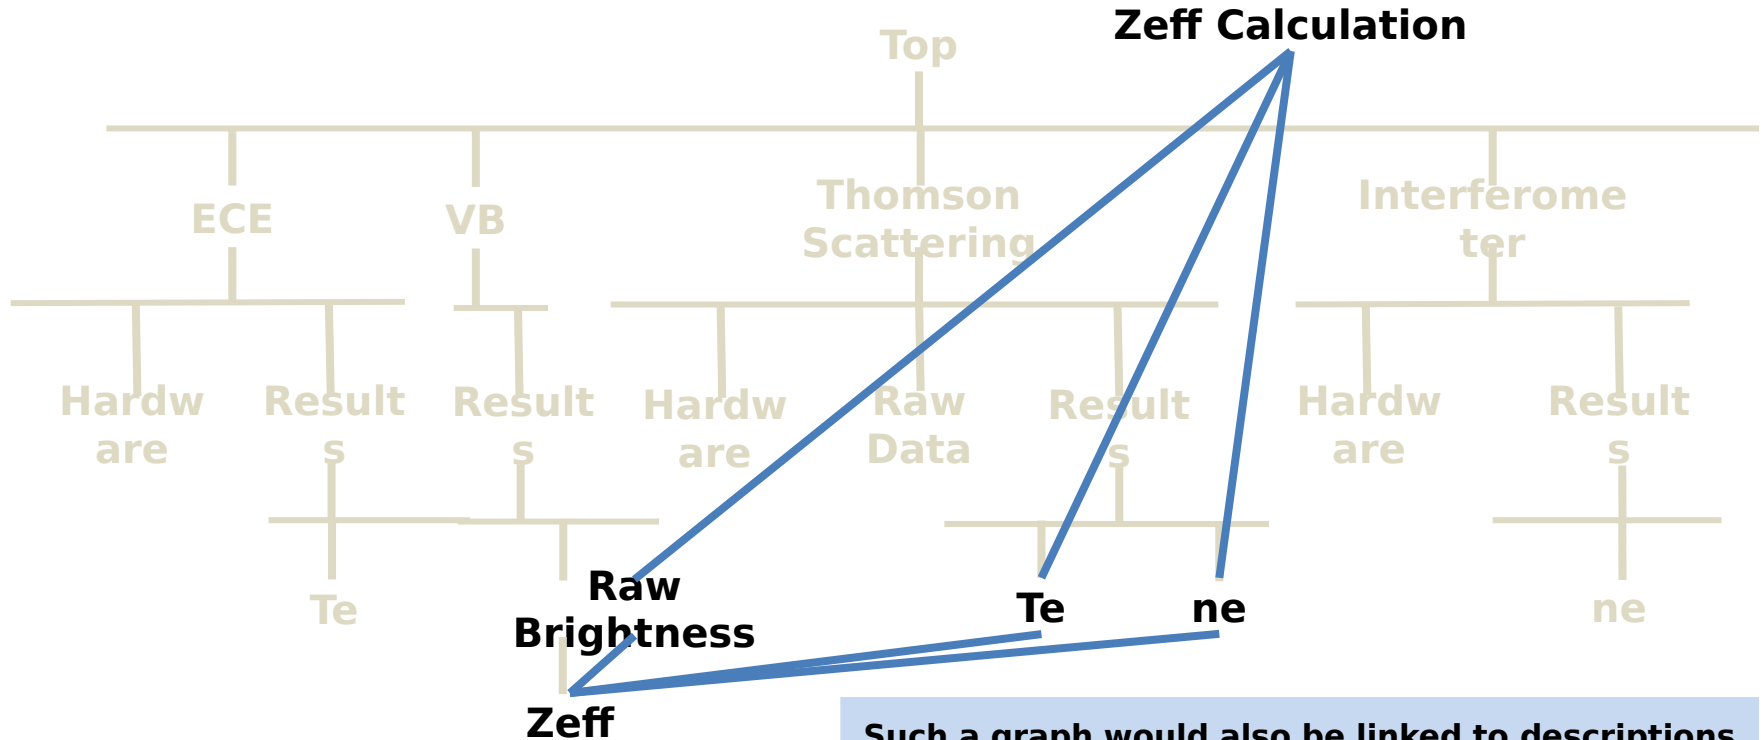
**The C-Mod Data System Has  $\sim 10^5$  Such Nodes With Significant Metadata For Each Node**

# Organization of Data - By Physics Parameter



Such a graph would also be linked to descriptions of experiments, annotation, etc.

# Organization of Data - By Data Provenance



Such a graph would also be linked to descriptions of analysis codes, annotation, etc.

# Approach

- Use graphs to describe relationships between data
- Schema defined using schema.JSON











## – Nodes






- Who, what, when, history
- List of properties appropriate to their type
- [URI to have objects stored in other systems]
  - Protocol://location/specifiers

◦ GUID

## – Edges

- Who, what, when, history
- Allowable SRCs, Destinations
- Properties (if needed)

Name ?	Color	SuperClasses ?	Name ?	Color	SuperClasses ?
Annotation		_NDMObject, V	MiniProposal		_ExternalReference
CModRunDay		_RunDay	PDFReference		_ExternalReference
CModShot		_Shot	Person		V
DropboxImage		_ExternalReference	Topic		_NDMObject, V
MDSplus		_ExternalReference	V		

Name ?	Color	SuperClasses ?
CommentsOn		_NDMObject, E
E		
HasImage		_NDMObject, E
HasTopic		_NDMObject, E
RunOn		_NDMObject, E

# Approach

- Use graphs to describe relationships between data

- Schema defined using schema.JSON

- Nodes

- o Who, what, when, history

- o List of properties appropriate to their type

- o [URI to have objects stored in other systems]

- Protocol://location/specifiers

- o GUID

- Edges

- o Type

- o Allowable SRCs, Destinations

- o Properties (if needed)

```
$ cat PDFReference.json
{
  "id": "http://ndm.mit.edu/ndm/schemas/PDFReference",
  "$schema": "http://json-schema.org/draft-06/schema#",
  "title": "PDFReference",
  "description": "Link to a web accessible PDF",
  "definitions": {},
  "type": "object",
  "allof": [{ "$ref": "_ExternalReference" } ],
  "properties": {
    "name": { "type": "STRING" },
    "type": {
      "type": "STRING",
      "defaultvalue": "PDF"
    }
  },
  "required": ["name"],
  "metadata": {
    "schemaMetadata": {
      "title" : "{{@class}} {{$name}}",
      "brief" : [],
      "body" : ["URI"],
      "links" : []
    }
  }
}
```



## Approach

---

- Graph database - OrientDB
- Javascript SPA
- VUE.js - web frontend xxx
- Auth0
- Docker-compose (microservices)
  - DB server
  - Authenticator Proxy
  - Web Server
  - Notification server
- PM2 - lightweight process orchestration

```
$ cat docker-compose.yml  
version: '3'
```

```
services:
```

```
  web:
```

```
    build:
```

```
      context: './Client'
```

```
      args:
```

```
        - dbname=${ORIENT_DBNAME}
```

```
        - dropboxkey=${DROPBOX_KEY}
```

```
      ports:
```

```
        - 80:80
```

```
      depends_on:
```

```
        - proxy
```

```
  proxy:
```


```
    build: './Proxy'
```

```
    volumes:
```


```
      - ./Proxy/db:/usr/src/app/db
```

```
...
```

# Proposals, Run Days, Shots and Annotations



# Proposals, Run Days, Shots and Annotations



# Proposals, Run Days, Shots and Annotations

-  CModRunDay
-  MiniProposal
-  CModShot
-  Annotation
-  HasProposal
-  RunOn
-  CommentsOn

## Abstract

This experiment intends to explore the feasibility of operation in partially, pronounced and/or full detachment while maintaining a high-confinement,  $H_{98} \sim 1$ , I-mode pedestaland core. The goal is provide scoping and demonstration of a possible mixed low-Z seeding approach using Ne/N<sub>2</sub> which can then be followed up with demonstrations over a wider range of I-mode plasmas.

## CoAuthors

J.W. Hughes, A. Hubbard, B. Mumgaard, D. Brunner, B. LaBombard, J. Terry, A.Q. Kuang, S. Wolfe, I. Hutchinson, J. Canik (ORNL), C. Thielier (CRPP), B. Lipschultz (U. of York)

**submitDate**

2019-04-29 16:41:53



# Proposals, Run Days, Shots and Annotations

← → ↻ 🏠 ⓘ localhost:8080/miniprops/63:1 🔍 ☆ 📧 📺 🟢 | 👤 ⋮

## NDM

### Exploring compatibility of detachment with I-mode plasmas

By **Reinke, Matt**, J.W. Hughes, A. Hubbard, B. Mumgaard, D. Brunner, B. LaBombard, J. Terry, A.Q. Kuang, S. Wolfe, I. Hutchinson, J. Canik (ORNL), C. Thieler (CRPP), B. Lipschultz (U. of York)

This experiment intends to explore the feasibility of operation in partially, pronounced and/or full detachment while maintaining a high-confinement,  $H_{98} \sim 1$ , I-mode pedestal and core. The goal is provide scoping and demonstration of a possible mixed low-Z seeding approach using Ne/N<sub>2</sub> which can then be followed up with demonstrations over a wider range of I-mode plasmas.

### Assessing time-evolving particle transport across I/H transitions

By **Reinke, Matt**, A. Loarte (ITER), M. Cerretti (U. of York), J.W. Hughes, J.E. Rice, A. Hubbard, J. Walk, J. Terry, M. Chilenski, E. Edlund (PPPL), B. Mumgaard, S. Wolfe.

This experiment looks to investigate particle transport across the transition from an I-mode to an ELM-free H-mode, examining the time evolution of both impurity and main ion densities. The primary goal is a search for a solid existence proof, at least transiently, of radially outward high-Z

# Proposals, Run Days, Shots and Annotations

The screenshot shows a web browser window with the address bar displaying 'localhost:8080/miniprops/63:1'. The page has a dark blue header with the text 'NDM'. Below the header, there are two yellow proposal cards. The first card is titled 'Exploring compatibility of detachment with I-mode plasmas' and lists authors: 'By Reinke, Matt, J.W. Hughes, A. Hubbard, B. Mumgaard, D. Brunner, B. LaBombard, J. Terry, A.Q. Kuang, S. Wolfe, I. Hutchinson, J. Canik (ORNL), C. Thieler (CRPP), B. Lipschultz (U. of York)'. The text below the title describes the experiment's goal: 'This experiment intends to explore the feasibility of operation in partially, pronounced and/or full detachment while maintaining a high-confinement, H98 ~ 1, I-mode pedestal and core. The goal is provide scoping and demonstration of a possible mixed low-Z seeding approach using Ne/N2 which can then be followed up with demonstrations over a wider range of I-mode plasmas.' The second card is titled 'Assessing time-evolving particle transport across I/H transitions' and lists authors: 'By Reinke, Matt, A. Loarte (ITER), J. Walk, J. Terry, M. Chilenski, E. Edl...'. A small menu is open over this card with options 'Show MiniProposal' and 'Show Runs'. The text below the title describes the experiment: 'This experiment looks to investigate particle transport across the transition from an I mode to an ELM-free H-mode, examining the time evolution of both impurity and main ion densities. The primary goal is a search for a solid existence proof, at least transiently, of radially outward high-Z'.

← → ↻ 🏠 ⓘ localhost:8080/miniprops/63:1 🔍 ☆ 📧 📅 📌 👤 ⋮

## NDM

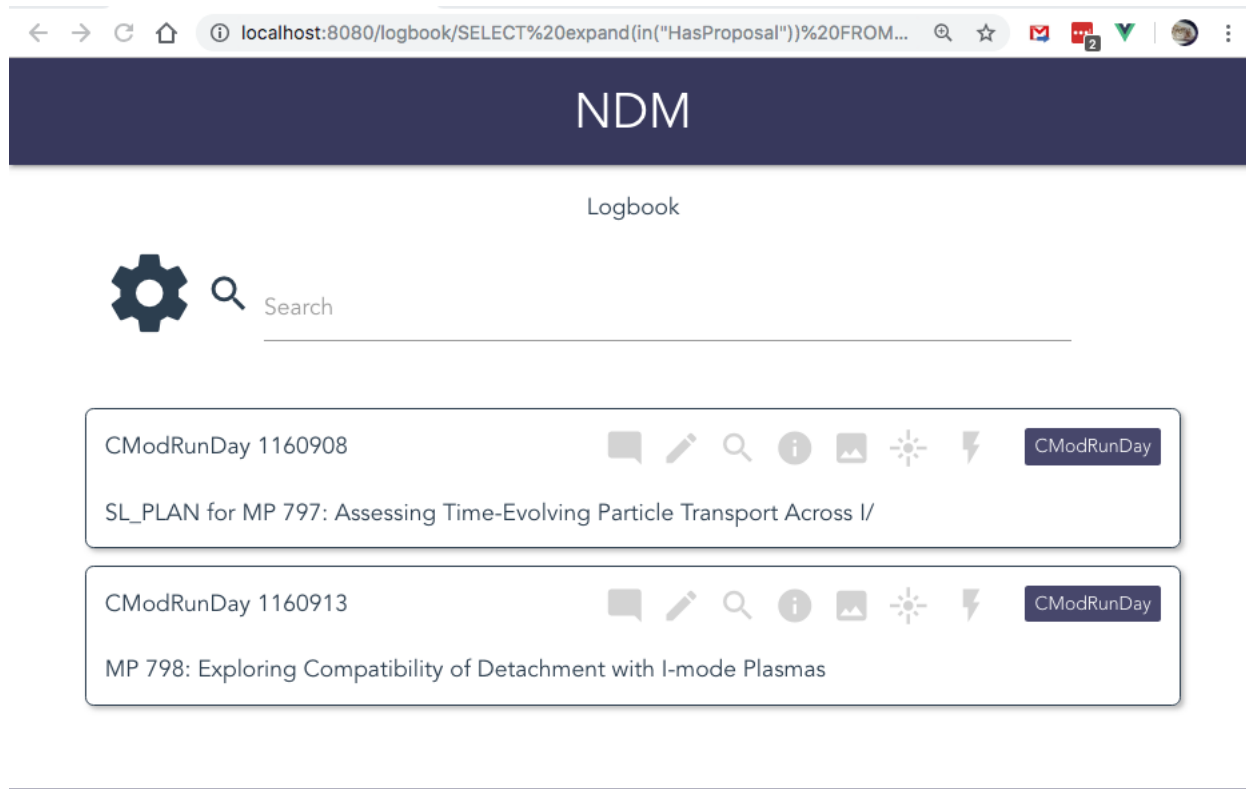
**Exploring compatibility of detachment with I-mode plasmas**  
By **Reinke, Matt**, J.W. Hughes, A. Hubbard, B. Mumgaard, D. Brunner, B. LaBombard, J. Terry, A.Q. Kuang, S. Wolfe, I. Hutchinson, J. Canik (ORNL), C. Thieler (CRPP), B. Lipschultz (U. of York)

This experiment intends to explore the feasibility of operation in partially, pronounced and/or full detachment while maintaining a high-confinement, H98 ~ 1, I-mode pedestal and core. The goal is provide scoping and demonstration of a possible mixed low-Z seeding approach using Ne/N2 which can then be followed up with demonstrations over a wider range of I-mode plasmas.

**Assessing time-evolving particle transport across I/H transitions**  
By **Reinke, Matt**, A. Loarte (ITER), J. Walk, J. Terry, M. Chilenski, E. Edl...  
Show MiniProposal  
Show Runs

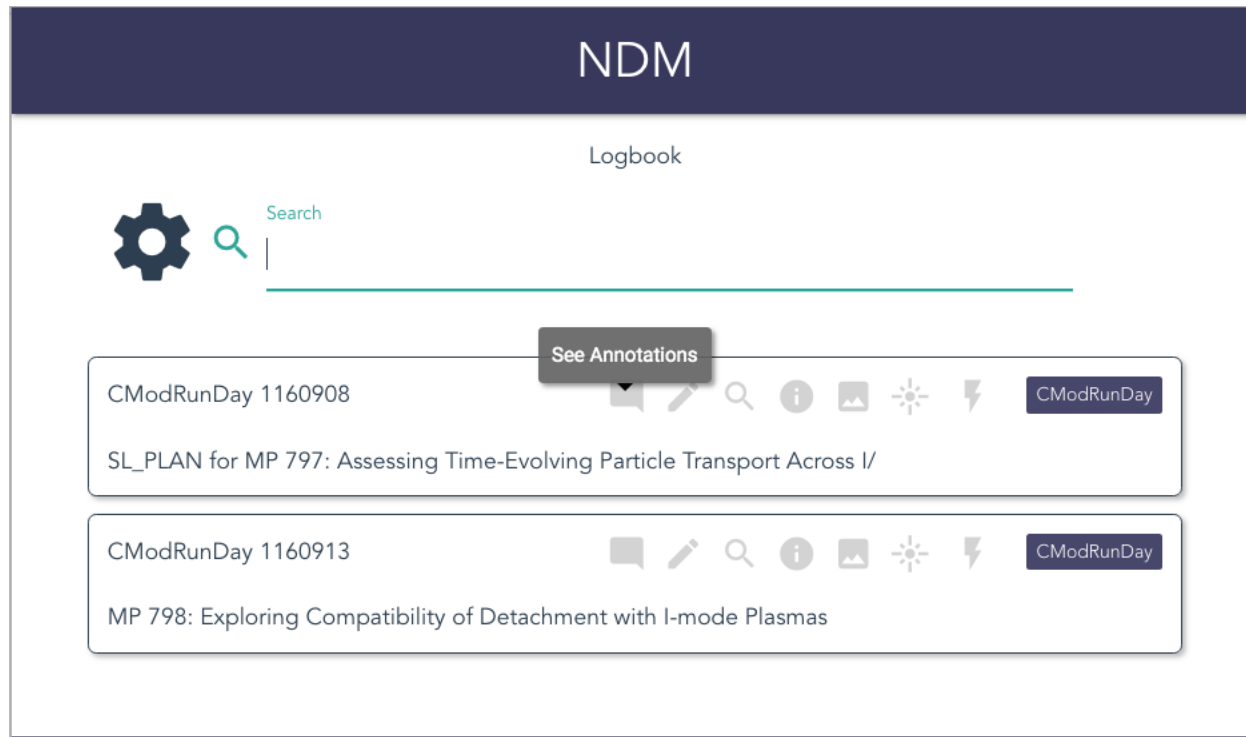
This experiment looks to investigate particle transport across the transition from an I mode to an ELM-free H-mode, examining the time evolution of both impurity and main ion densities. The primary goal is a search for a solid existence proof, at least transiently, of radially outward high-Z

# Proposals, Run Days, Shots and Annotations





# Proposals, Run Days, Shots and Annotations

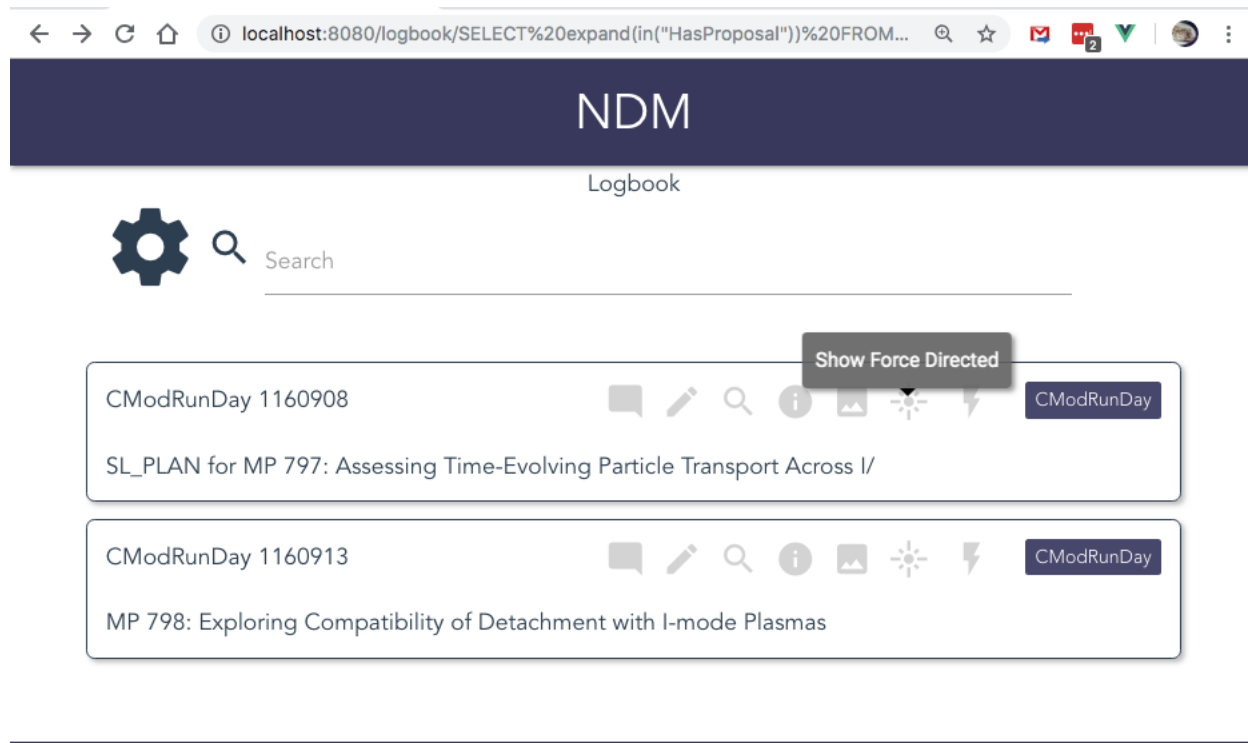


# Proposals, Run Days, Shots and Annotations


The screenshot shows a web browser window with the URL `localhost:8080/logbook/SELECT%20expand(in("HasProposal"))%20FROM...`. The application has a dark blue header with the text "NDM" and "Logbook" below it. A search bar with a gear icon and the word "Search" is visible. Below the header, a modal window is open, displaying a log entry for "CModRunDay 1160908". The modal has a title bar with icons for chat, edit, search, info, image, settings, and a lightning bolt, along with a "CModRunDay" button. The log entry text is as follows:

SL\_PLAN for MP 797: Assessing Time-Evolving Particle Transport Across I/  
**Author:** Matt Reinke **Topic:** ' SL\_PLAN  
**Title:** SL\_PLAN for MP 797: Assessing Time-Evolving Particle Transport Across  
SL\_PLAN for MP 797: Assessing Time-Evolving Particle Transport Across I/H Tran:  
SL: M.L. Reinke  
PO: "You sending the Wolfe? That's all you had to say!"  
The purpose of this run to test a hypothesis about the possible formation of  
of an outward directed impurity flux in the edge transport barrier in cases of  
high-temp, opaque pedestal. Indications of this happening at I/H transitions d

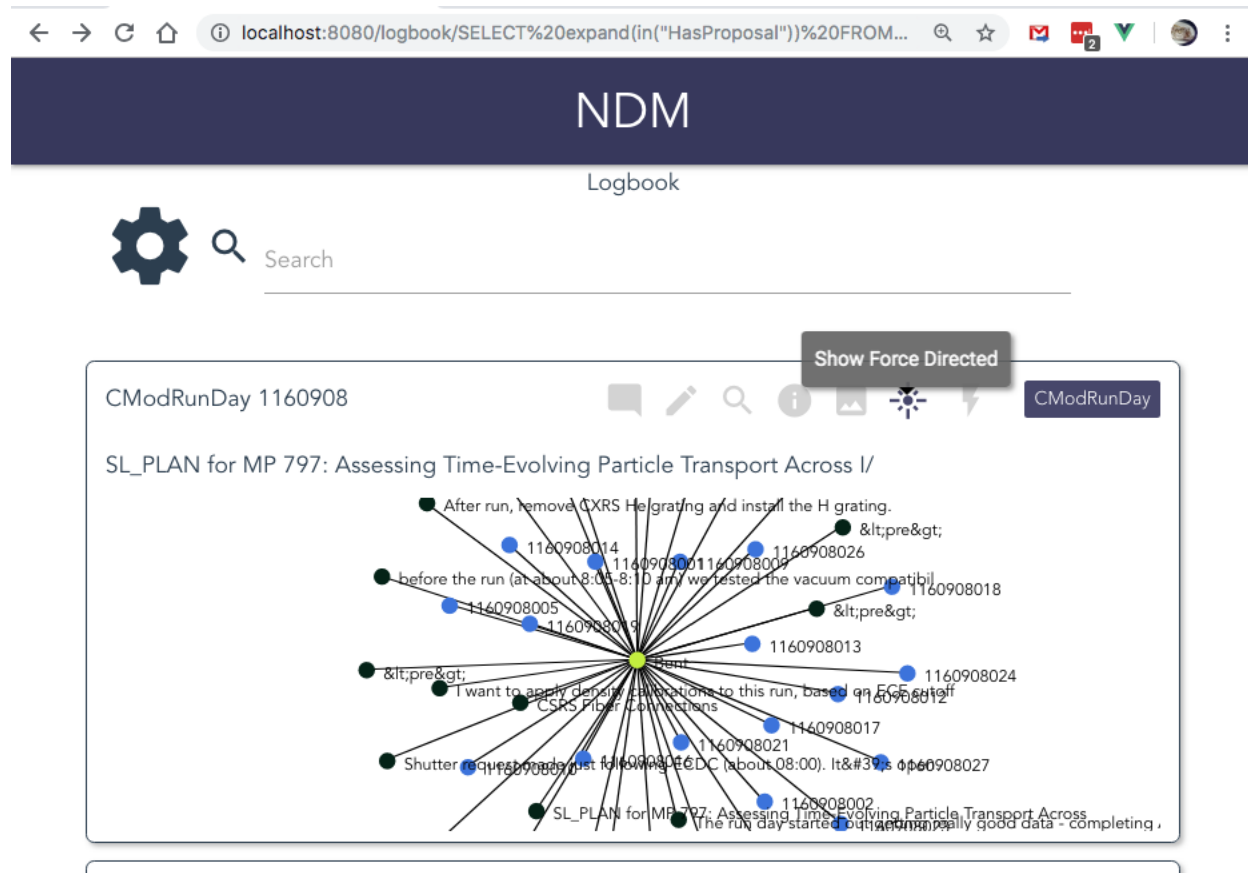
# Proposals, Run Days, Shots and Annotations



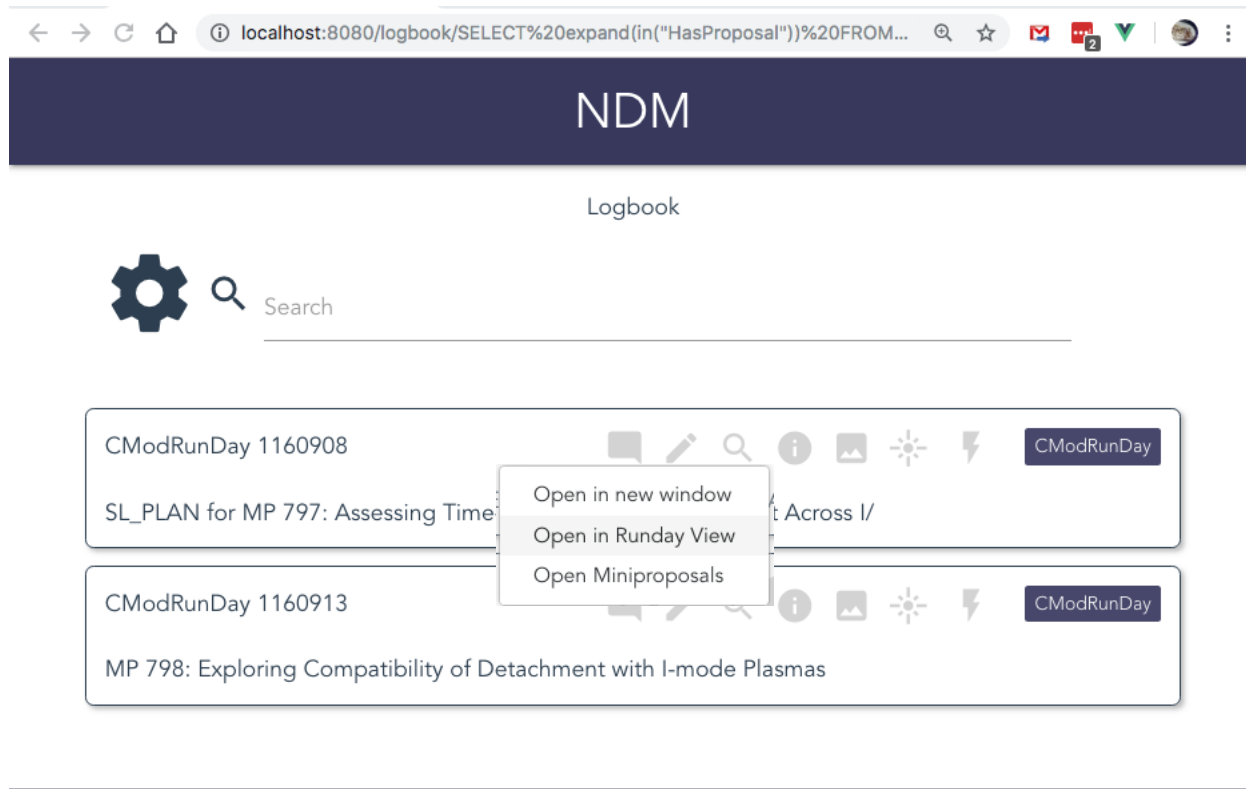
# Proposals, Run Days, Shots and Annotations



# Proposals, Run Days, Shots and Annotations



# Proposals, Run Days, Shots and Annotations



# Proposals, Run Days, Shots and Annotations

## CModRunDay 1160913

MP 798: Exploring Compatibility of Detachment with I-mode Plasmas

Shot - Monday Apr 29th 04:40 pm

Ramp Up Rate: Ramp Down Rate:  
Current Set Point Hold Time: Low Volatage Limit in Volts:

**Author:** Amanda Hubbard **Topic:** 'ECE

**Title:** For Matt's run (~5.4 T) I have changed GPC2 to 1.65 mm, 0 degrees. An

For Matt's run (~5.4 T) I have changed GPC2 to 1.65 mm, 0 degrees. And loaded xcal from 1160908001, which is like 1160908000. Leaving GPC set up for 8 T, since we'll be back at that on Wed. So use GPC2 and FRCECE for this run.

**Author:** Robert Mumgaard **Topic:** ' SPECTROSCOPY

**Title:** In preparation for the run I switched H-D from KTOP 4 via R2 4 to KTOP

In preparation for the run I switched H-D from KTOP 4 via R2 4 to KTOP 3 via R2 3 since Matt requested KTOP 4 be

**Author:** Robert Mumgaard **Topic:** ' SPECTROSCOPY

**Title:** Setting up the spectroscopy.

Setting up the spectroscopy.

CHROMEX is setup as 1150923:

A\_BOT 09, R1 13, NA, CHROMEX 01

A\_BOT 11, R1 06, NA, CHROMEX 02

A\_BOT 13, R1 08, NA, CHROMEX 03

A\_BOT 15, R1 01, NA, CHROMEX 04

K\_BOT 04, R1 17, XOVR 01, PHOTON2 B1

K\_BOT 05, R1 09, NA, CHROMEX 05

K\_BOT 06, R1 18, XOV R 02, PHOTON2 B2

K\_BOT 07, R1 02, NA, CHROMEX 06

K BOT 08, R1 12, XOVR 03, PHOTON2 B3

**⚡ MAKE SHOT**

 UPDATE SHOT

Shot Number

Ramp Up Rate

Ramp Down Rate

Current Set Point

Current Set Point Hold Time

Low Volatage Limit in Volts

High Voltage Limit in Volts

Normal 

A     T<sub>x</sub>



# Conclusions

---

- Start with simple user interfaces
  - Too much complexity and fanciness inhibits iteration
  - Complexity must be carried along
  - Each user facing function has to implement the shared UI
- Applications need to be customized to their tasks
  - General backend with application specific front-end / user interfaces
  - Sharing underlying data structures and APIs
    - facilitates the development of applications
    - Allows for mixing and traversing of information domains
  - Application specific code is needed to achieve needed functionality and usability
- Early users are critical
  - Fancy initial user Interfaces complicates this
- Authentication is difficult
  - Authorization is even harder
- Initial project (and funding) winding down, will continue development.
- We need this to integrate the disparate information about our research.

---

# END

---