MAX-PLANCK-INSTITUT
FÜR PLASMAPHYSIK

On Intensity, Relativity, Shift & Compression

– Video data archiving at W7-X

Wendelstein 7-X

S.Fischer[1], L.Vandevenne[2], P.Perusek[3], M.Grün[1], A.Holtz[1], M.Grahl[1], J.Pribošek[3], J.Alakuijala[2] and the W7-X team[1]

[1] Max Planck Institute for Plasma Physics, Greifswald, DE - [2] Google Research, Zurich, CH - [3] Cosylab, d. d., Control System Laboratory, Ljubljana, SI

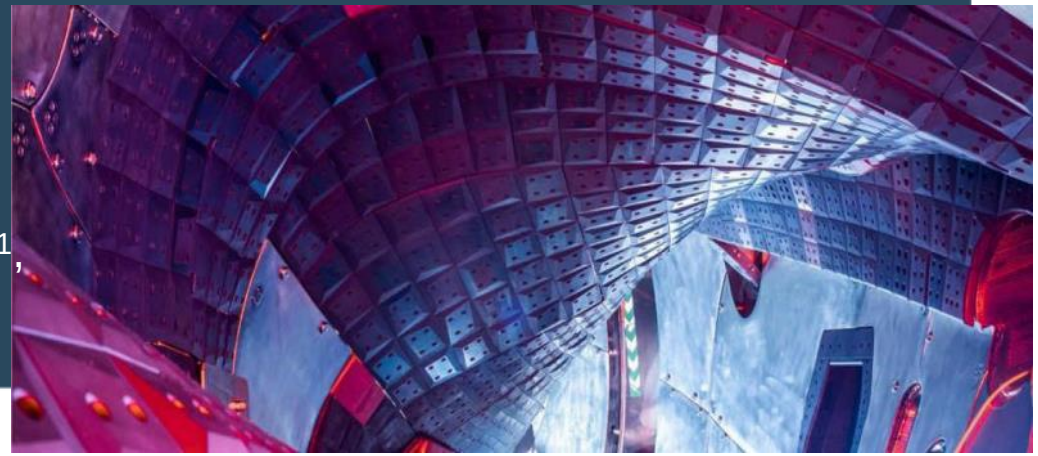

Credit: MPI for Plasma Physics

EUROfusion

This work has been carried out within the framework of the EUROfusion Consortium, funded by the European Union via the Euratom Research and Training Programme (Grant Agreement No 101052200 — EUROfusion). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them.

Wendelstein 7-X

Wendelstein 7-X

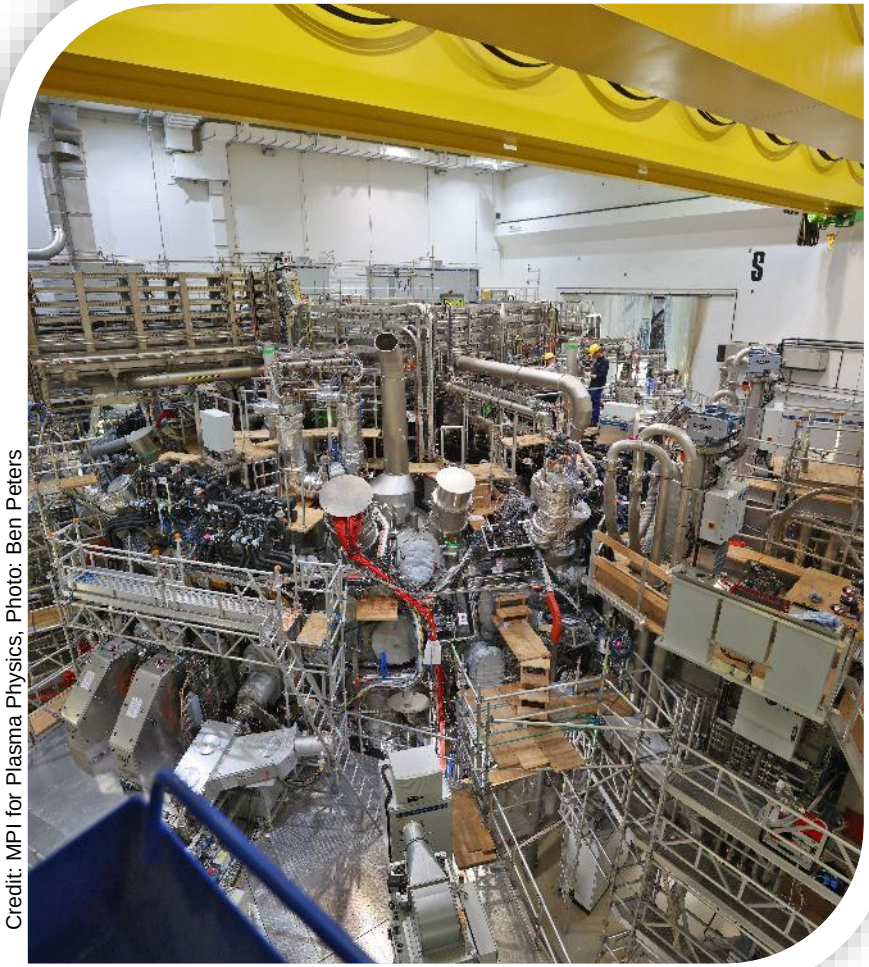

Credit: MPI for Plasma Physics, Photo: Ben Peters

Wendelstein 7-X

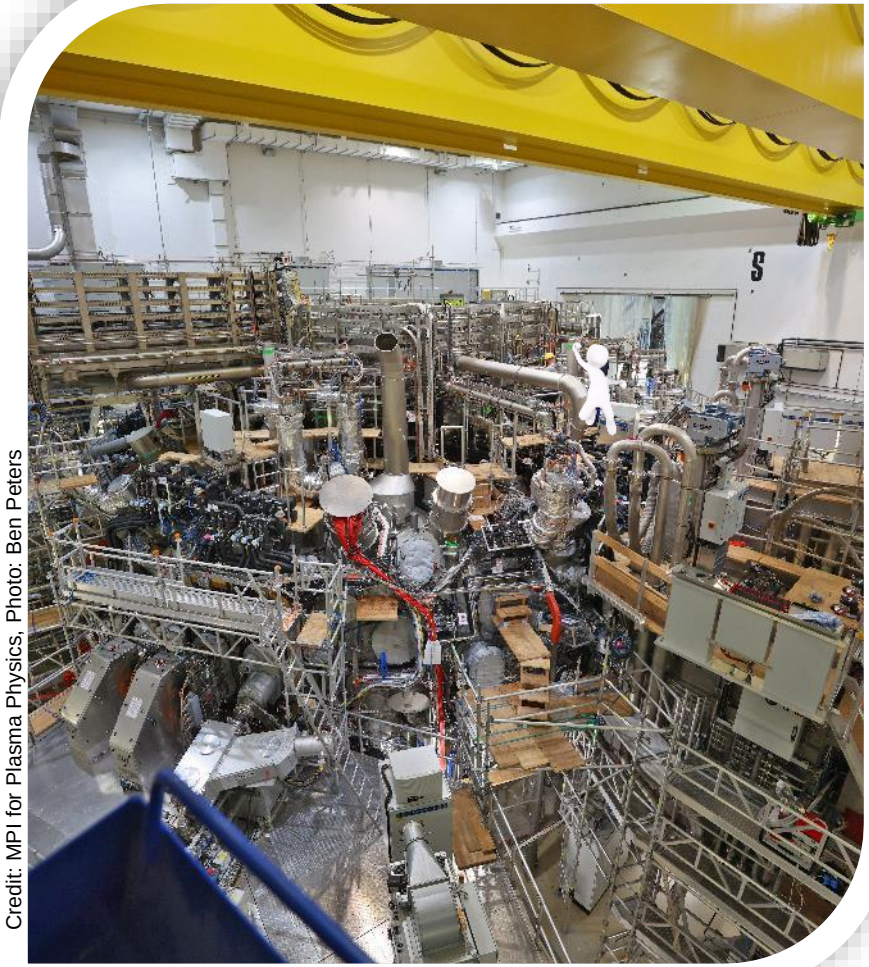

Credit: MPI for Plasma Physics, Photo: Ben Peters

Wendelstein 7-X

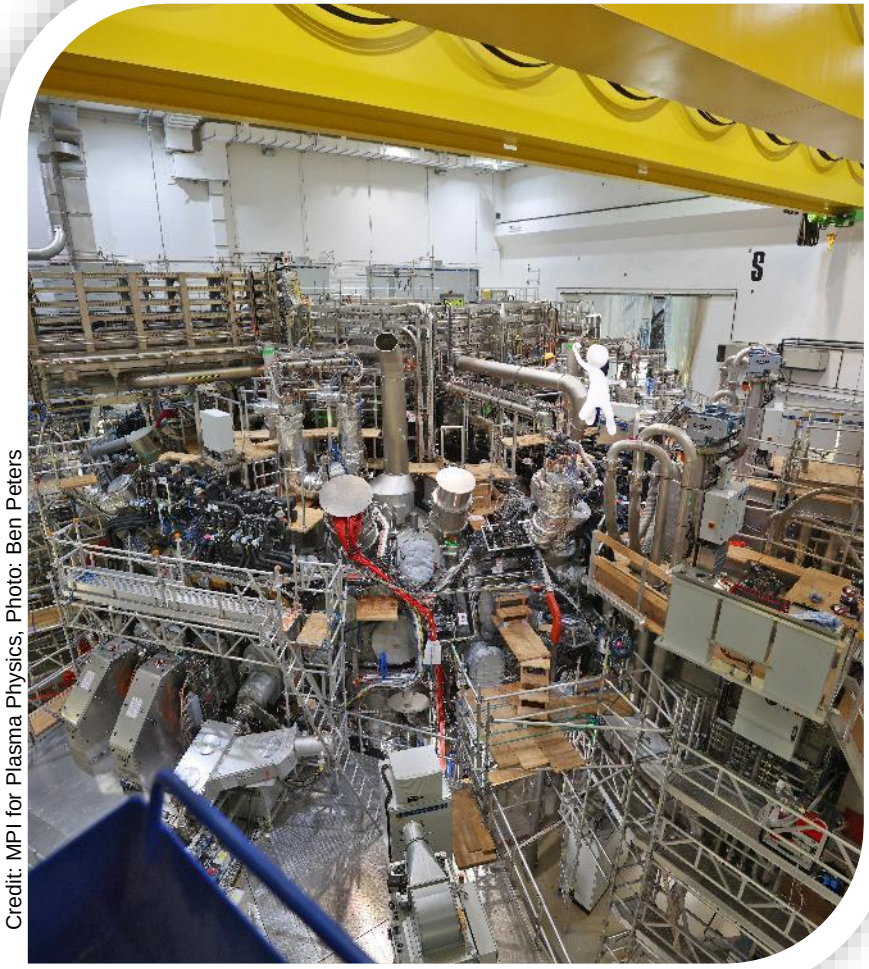

Credit: MPI for Plasma Physics, Photo: Ben Peters

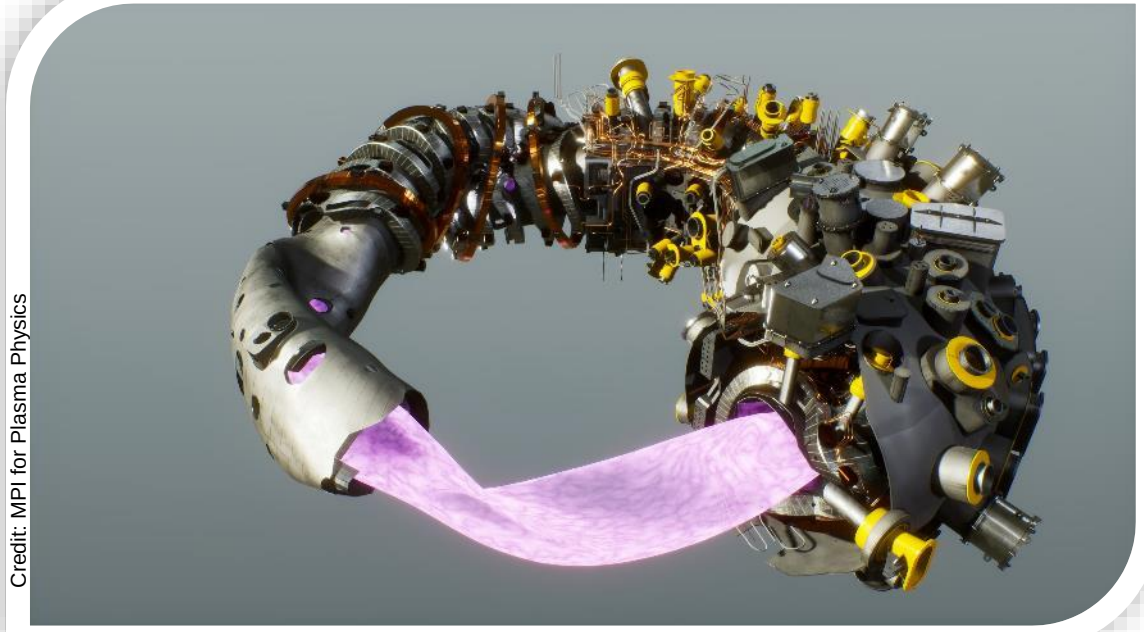

Credit: MPI for Plasma Physics

Wendelstein 7-X

Credit: MPI for Plasma Physics, Photo: Ben Peters

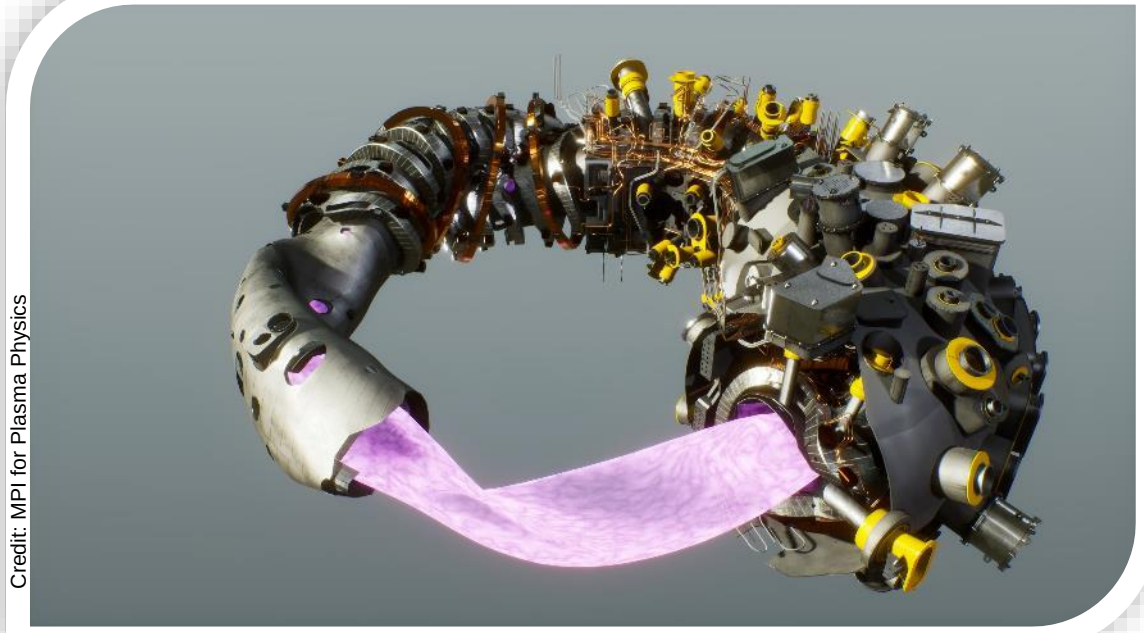

Credit: MPI for Plasma Physics

Wendelstein 7-X

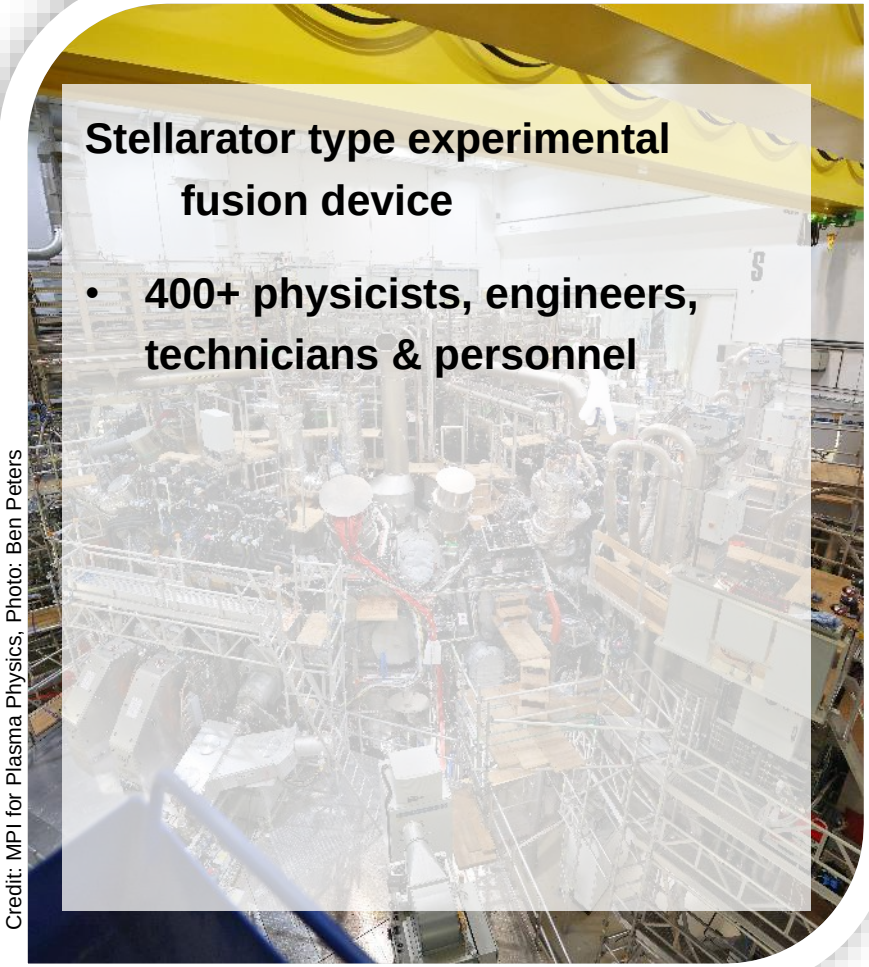

Credit: MPI for Plasma Physics, Photo: Ben Peters

Stellarator type experimental fusion device

- **400+ physicists, engineers, technicians & personnel**

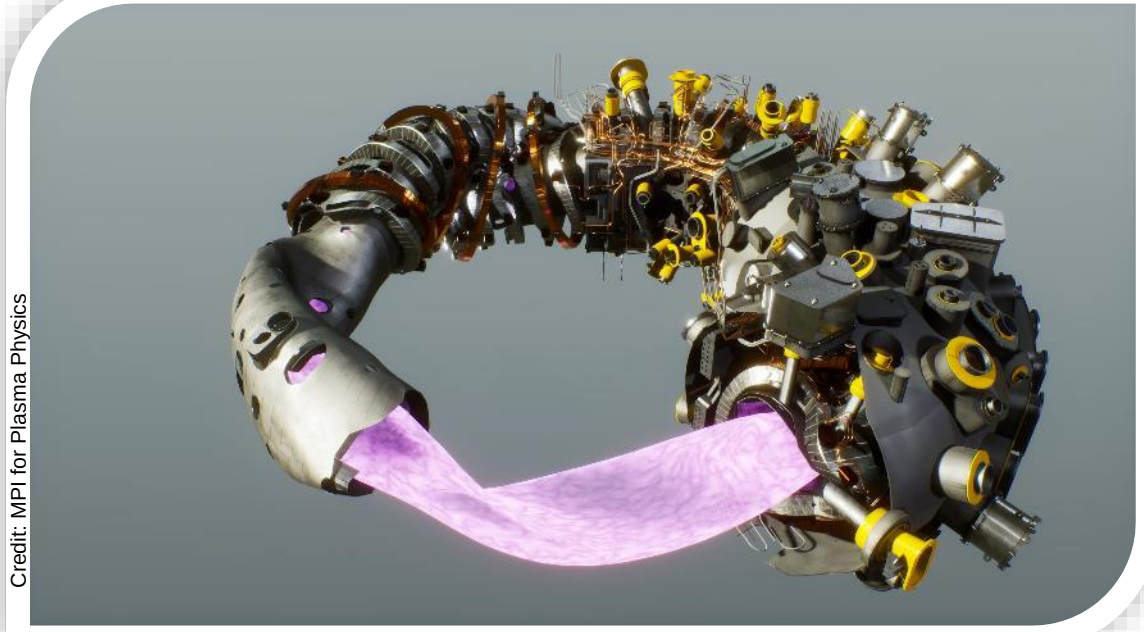

Credit: MPI for Plasma Physics

Wendelstein 7-X

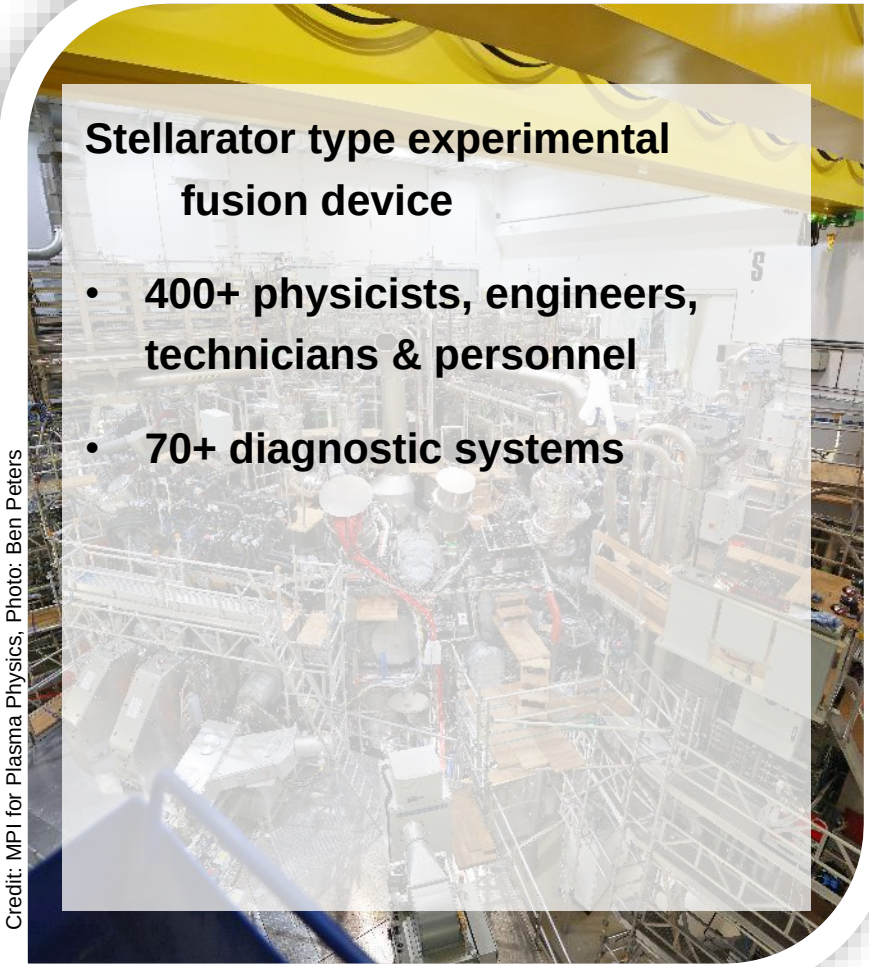

Credit: MPI for Plasma Physics, Photo: Ben Peters

Stellarator type experimental fusion device

- **400+ physicists, engineers, technicians & personnel**
- **70+ diagnostic systems**

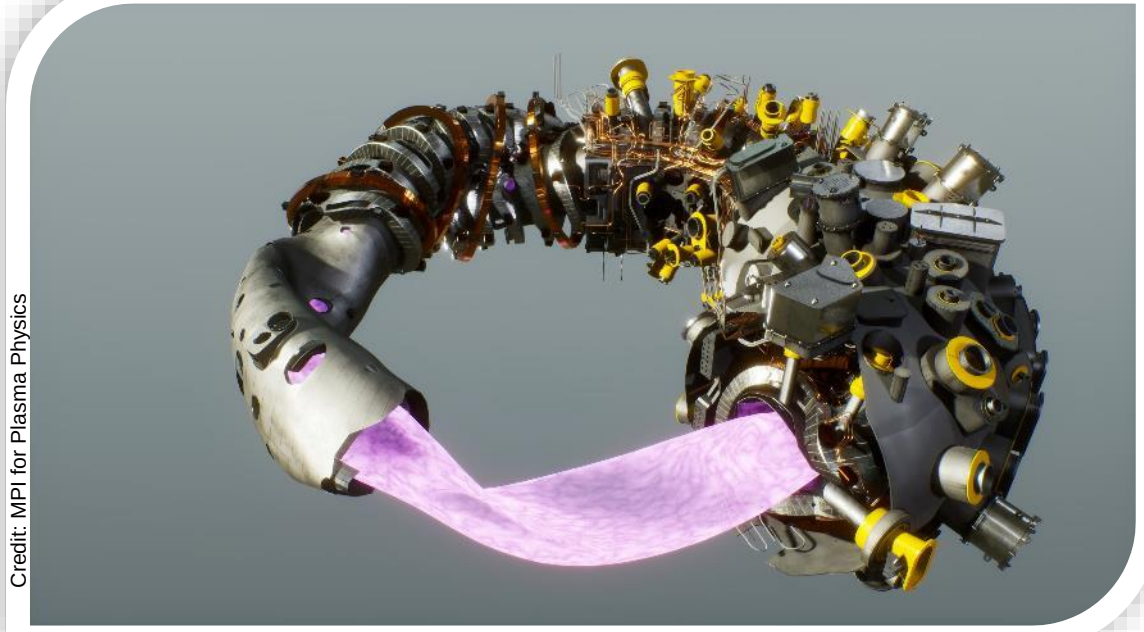

Credit: MPI for Plasma Physics

Wendelstein 7-X

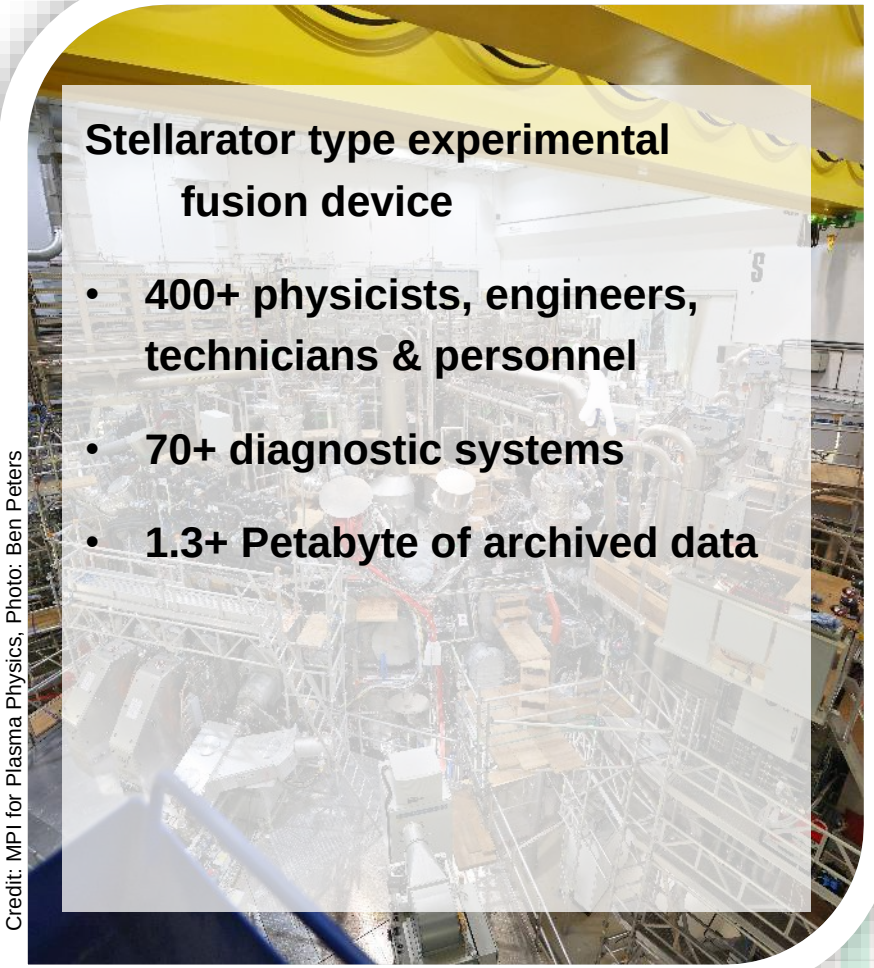

Credit: MPI for Plasma Physics, Photo: Ben Peters

Stellarator type experimental fusion device

- **400+ physicists, engineers, technicians & personnel**
- **70+ diagnostic systems**
- **1.3+ Petabyte of archived data**

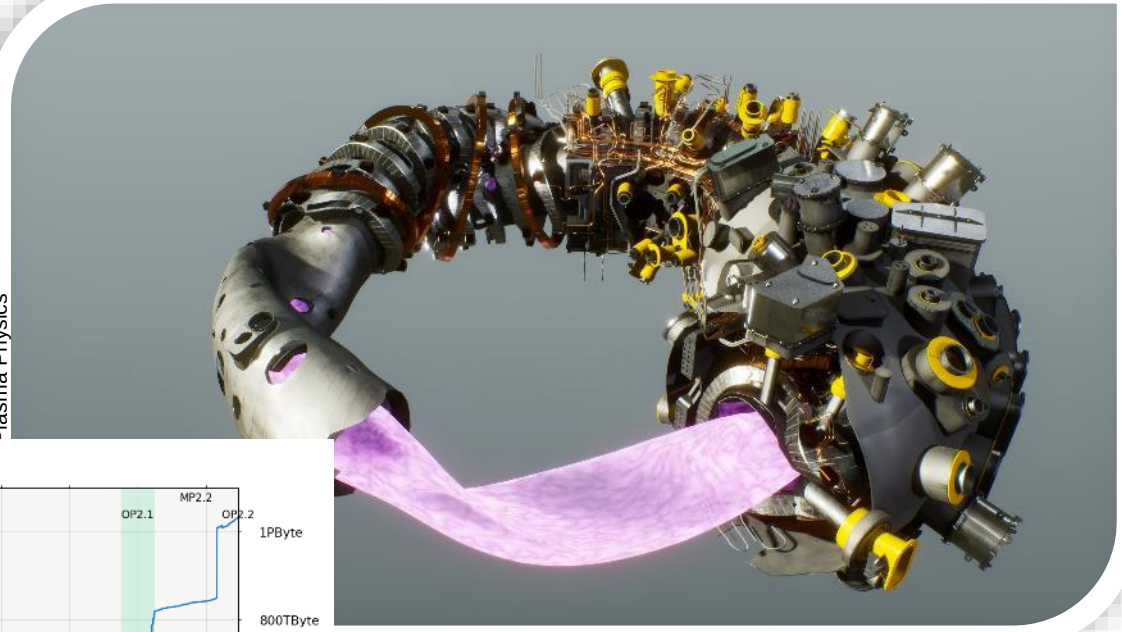

Plasma Physics

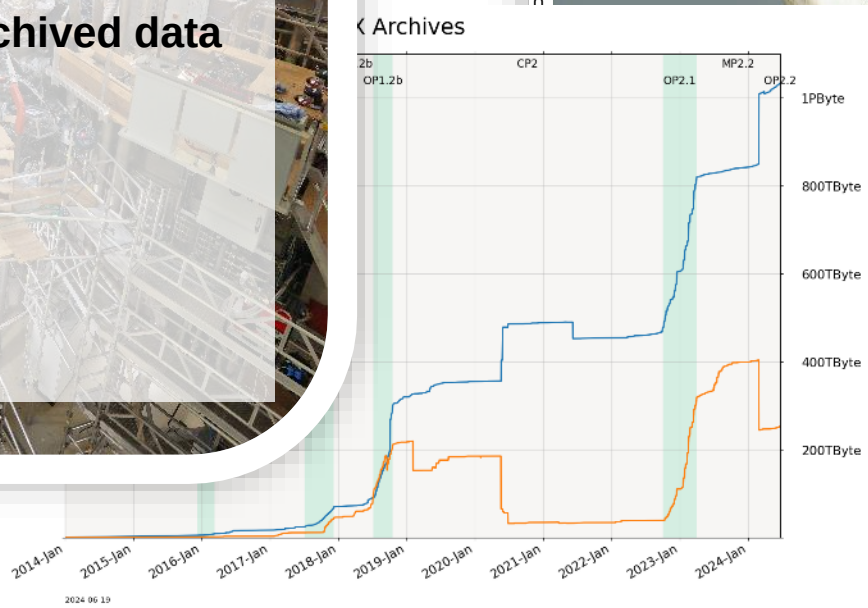

Wendelstein 7-X

Credit: MPI for Plasma Physics, Photo: Ben Peters

Stellarator type experimental fusion device

- **400+ physicists, engineers, technicians & personnel**
- **70+ diagnostic systems**
- **1.3+ Petabyte of archived data**
- **OP2.2 upcoming end of Q3**

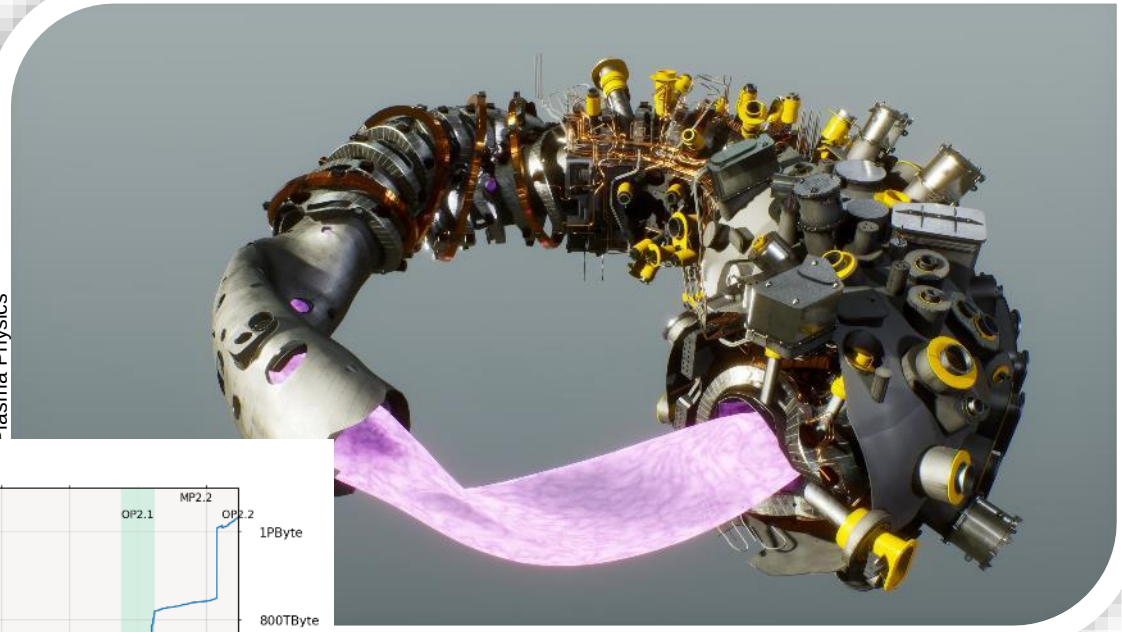

Plasma Physics

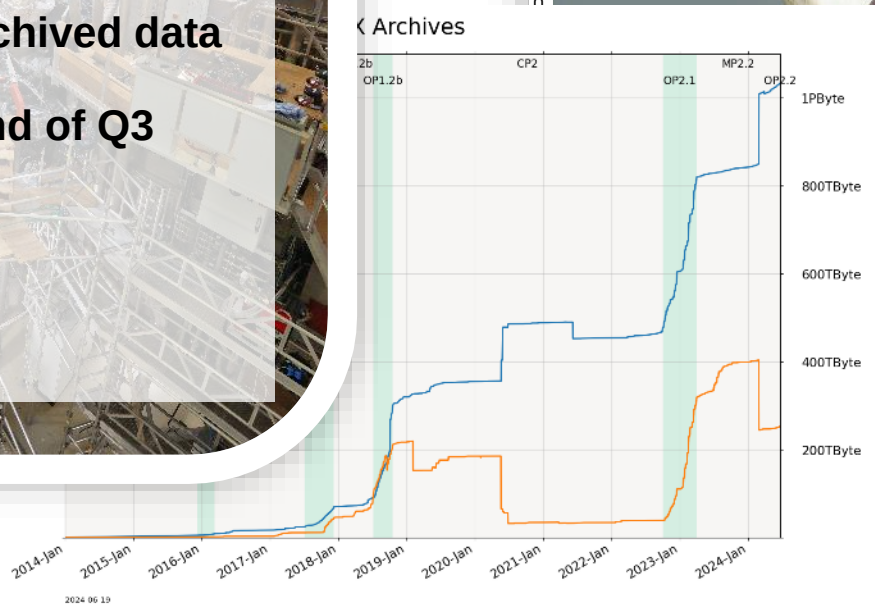

Wendelstein 7-X

Credit: MPI for Plasma Physics, Photo: Ben Peters

Stellarator type experimental fusion device

- **400+ physicists, engineers, technicians & personnel**
- **70+ diagnostic systems**
- **1.3+ Petabyte of archived data**
- **OP2.2 upcoming end of Q3**
- **1 campaign / year from now on: each ~65 experiments per day on ~40 days over 3 month**

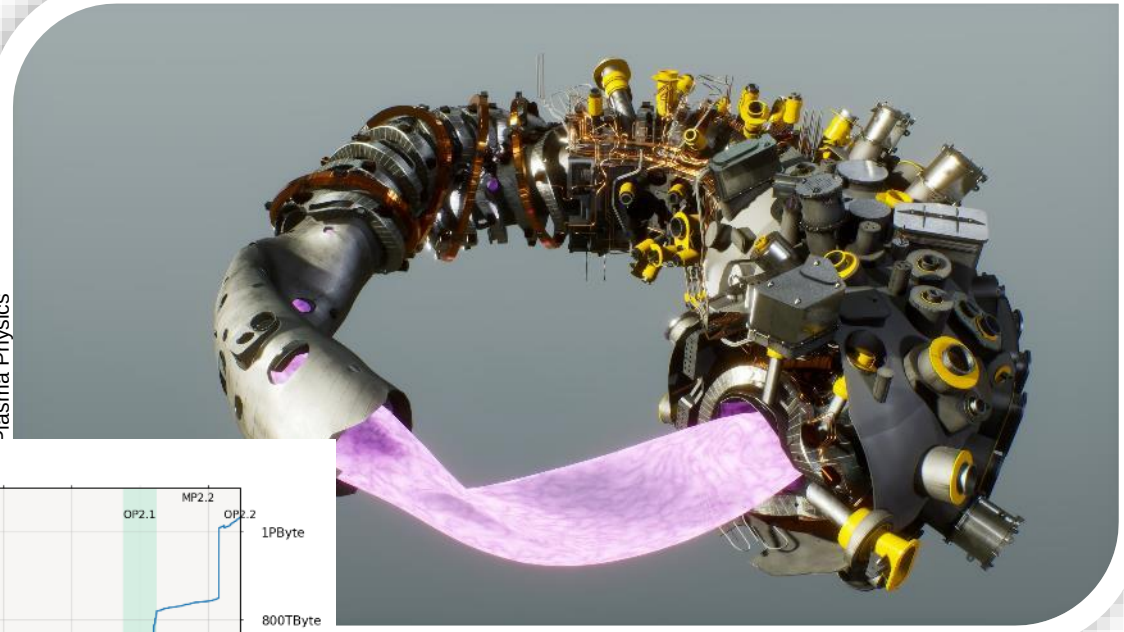

Plasma Physics

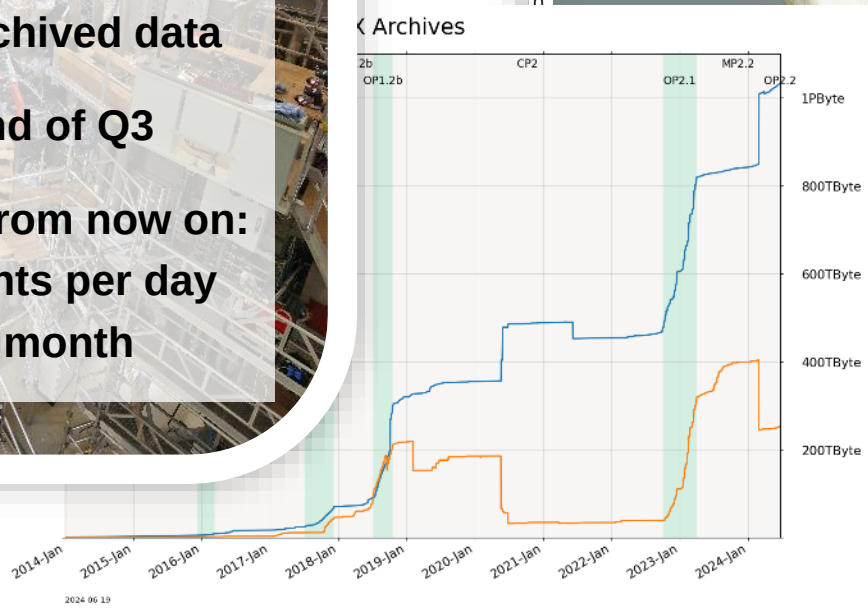

Wendelstein 7-X

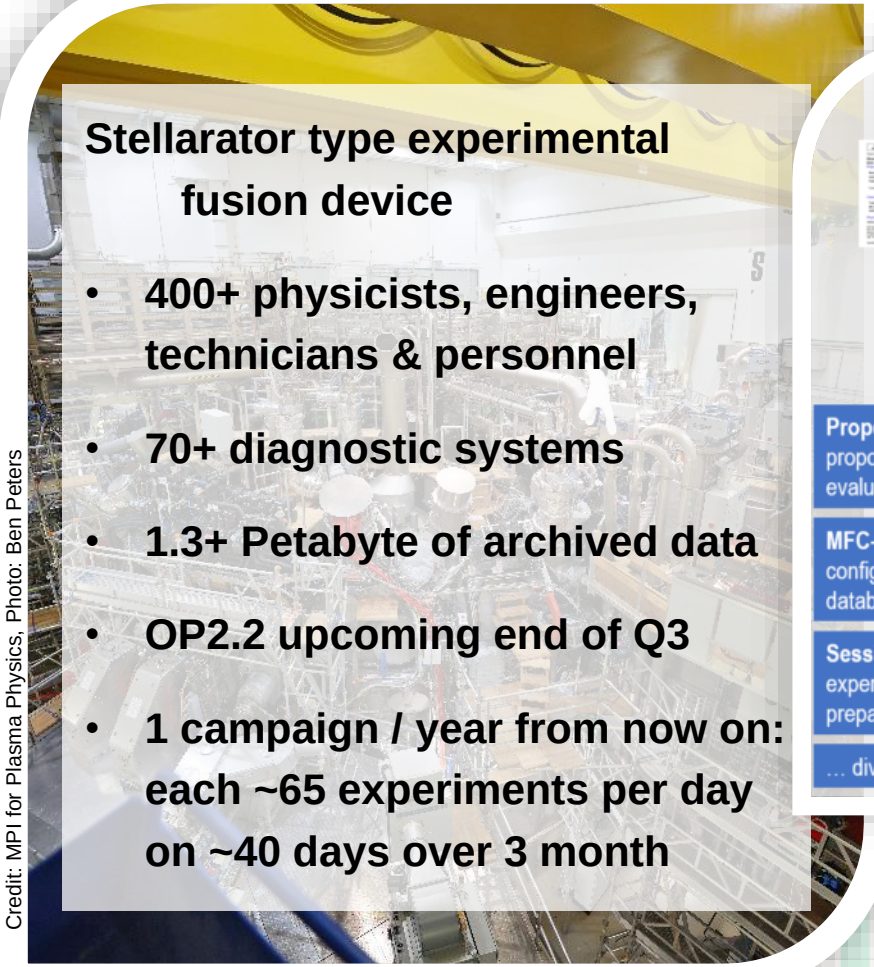

Credit: MPI for Plasma Physics, Photo: Ben Peters

Stellarator type experimental fusion device

- **400+ physicists, engineers, technicians & personnel**
- **70+ diagnostic systems**
- **1.3+ Petabyte of archived data**
- **OP2.2 upcoming end of Q3**
- **1 campaign / year from now on: each ~65 experiments per day on ~40 days over 3 month**

E5-CoDaC Software Development

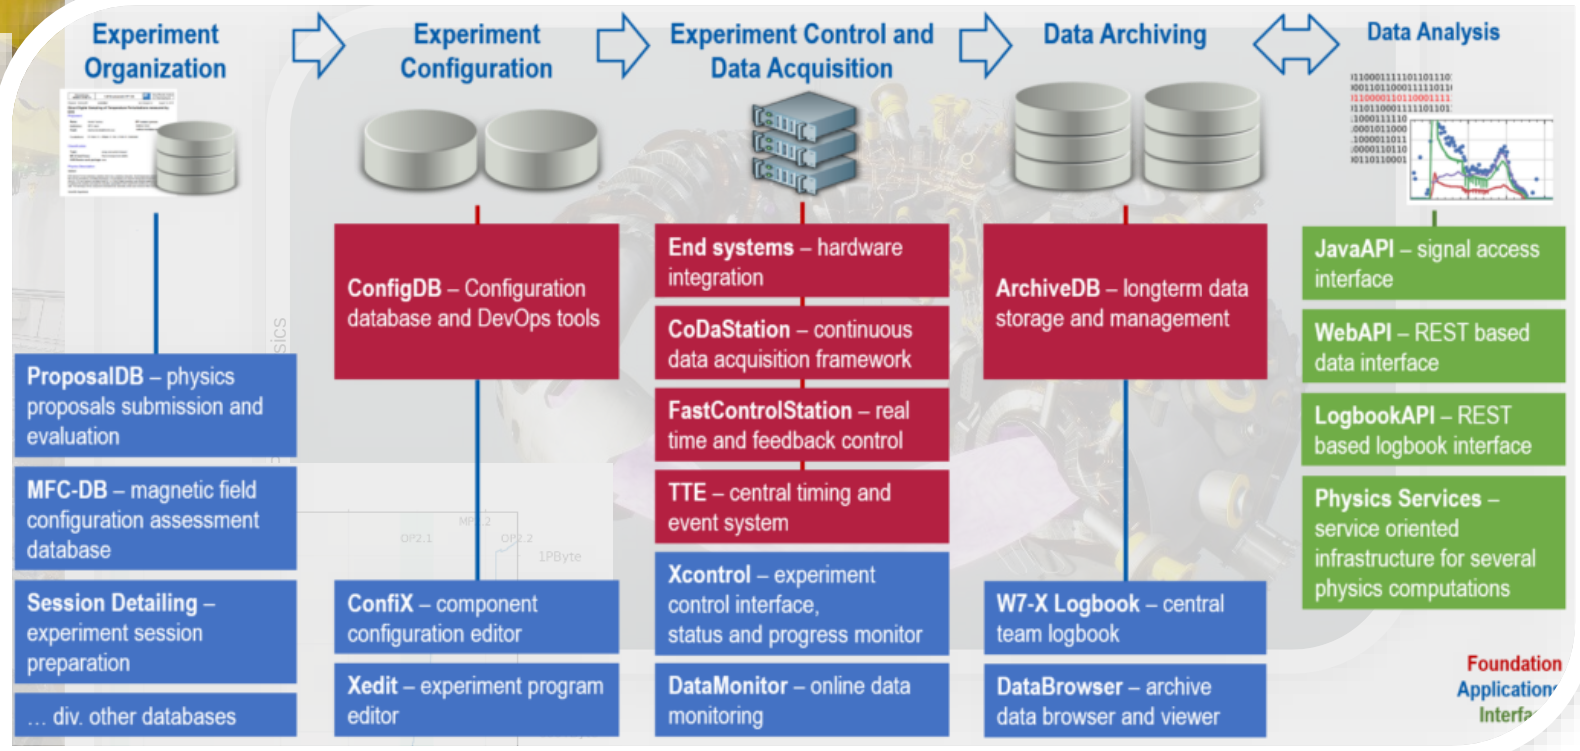

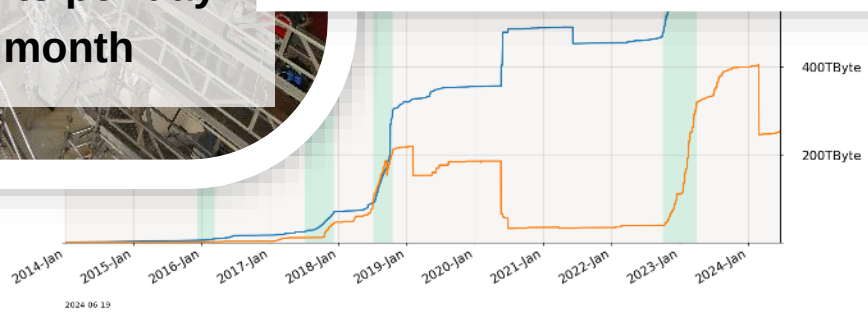

Wendelstein 7-X

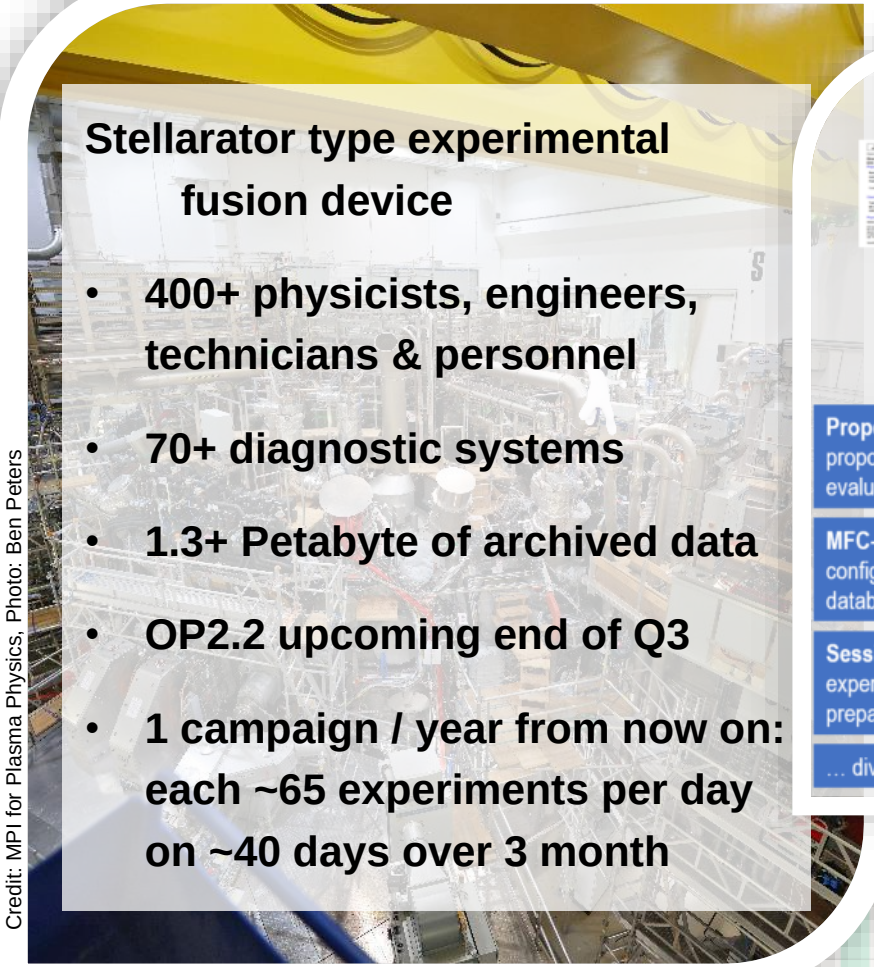

Credit: MPI for Plasma Physics, Photo: Ben Peters

Stellarator type experimental fusion device

- **400+ physicists, engineers, technicians & personnel**
- **70+ diagnostic systems**
- **1.3+ Petabyte of archived data**
- **OP2.2 upcoming end of Q3**
- **1 campaign / year from now on: each ~65 experiments per day on ~40 days over 3 month**

E5-CoDaC Software Development

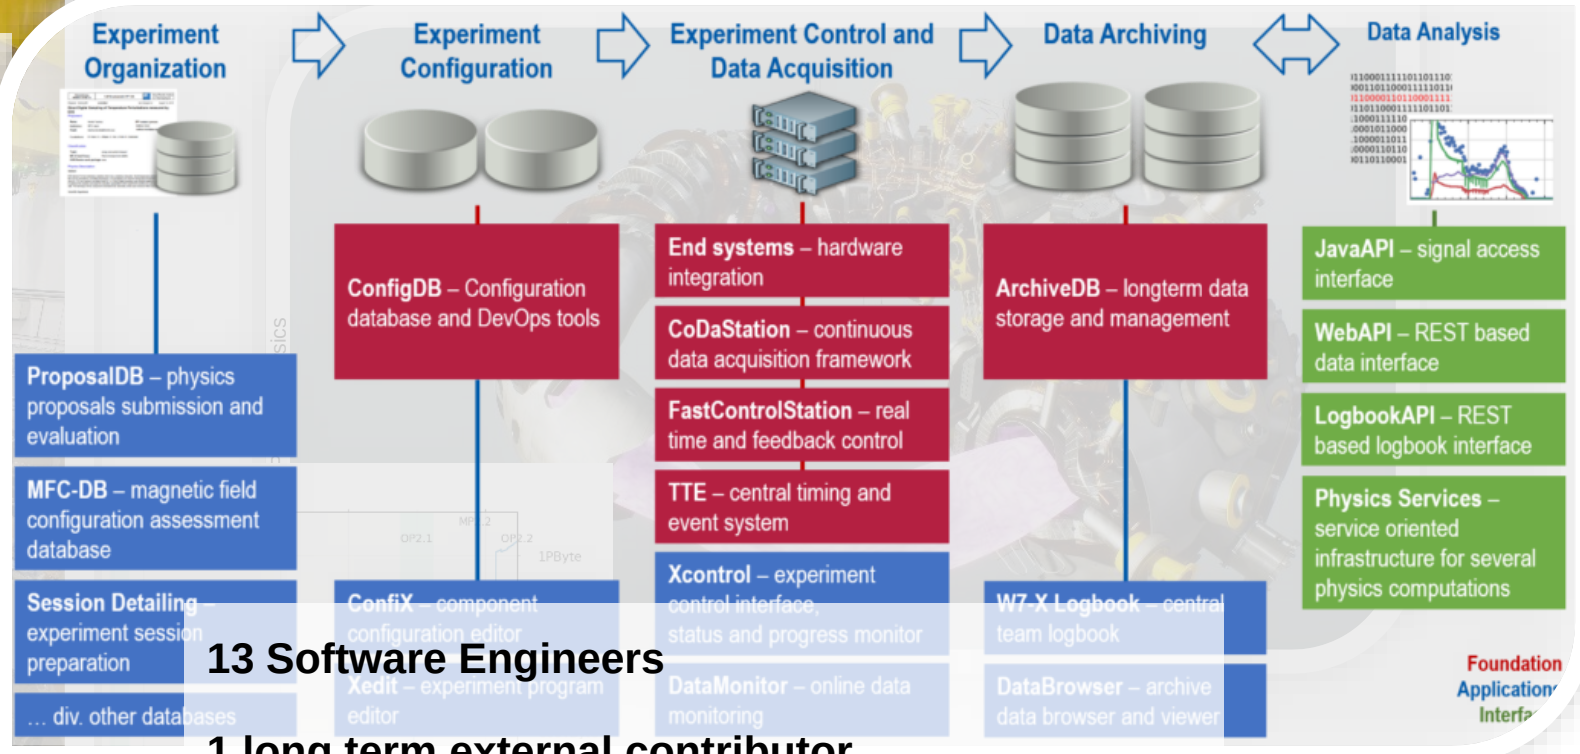

13 Software Engineers

1 long term external contributor

1 intern

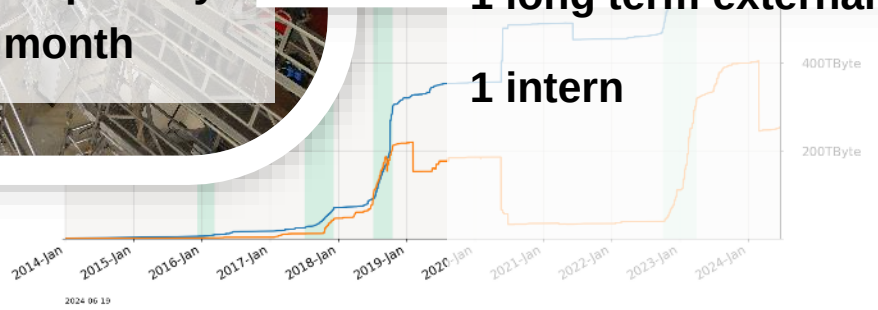

Wendelstein 7-X

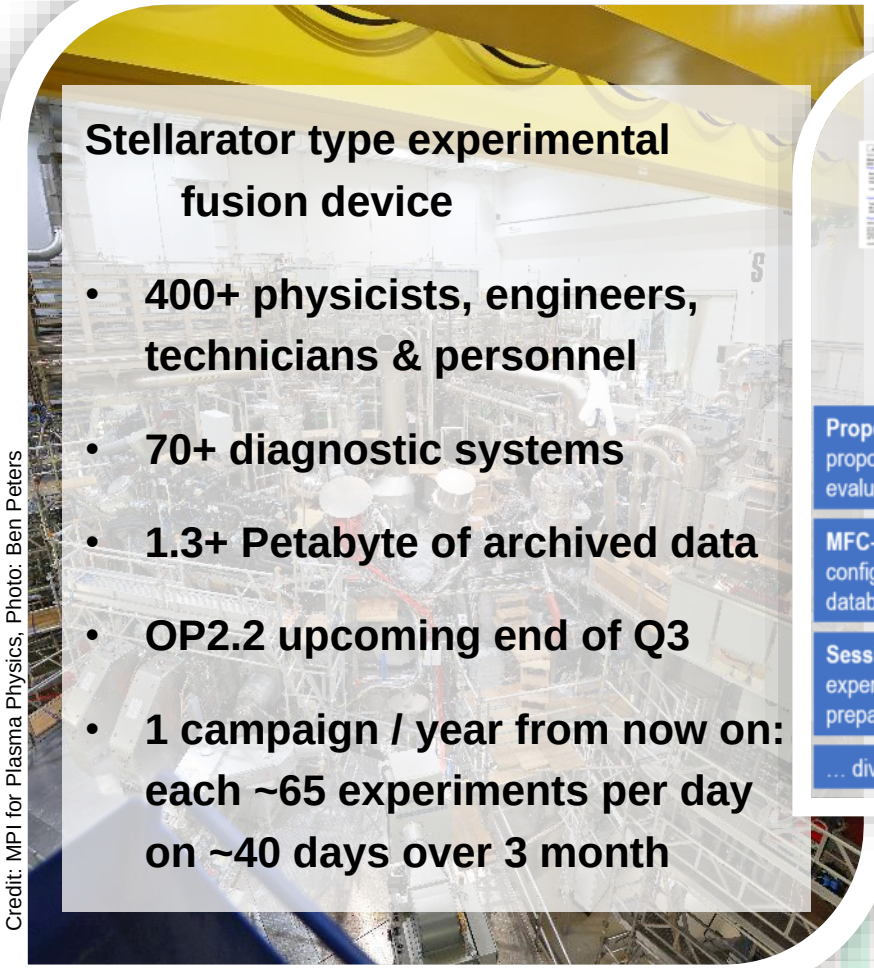

Credit: MPI for Plasma Physics, Photo: Ben Peters

Stellarator type experimental fusion device

- **400+ physicists, engineers, technicians & personnel**
- **70+ diagnostic systems**
- **1.3+ Petabyte of archived data**
- **OP2.2 upcoming end of Q3**
- **1 campaign / year from now on: each ~65 experiments per day on ~40 days over 3 month**

E5-CoDaC Software Development

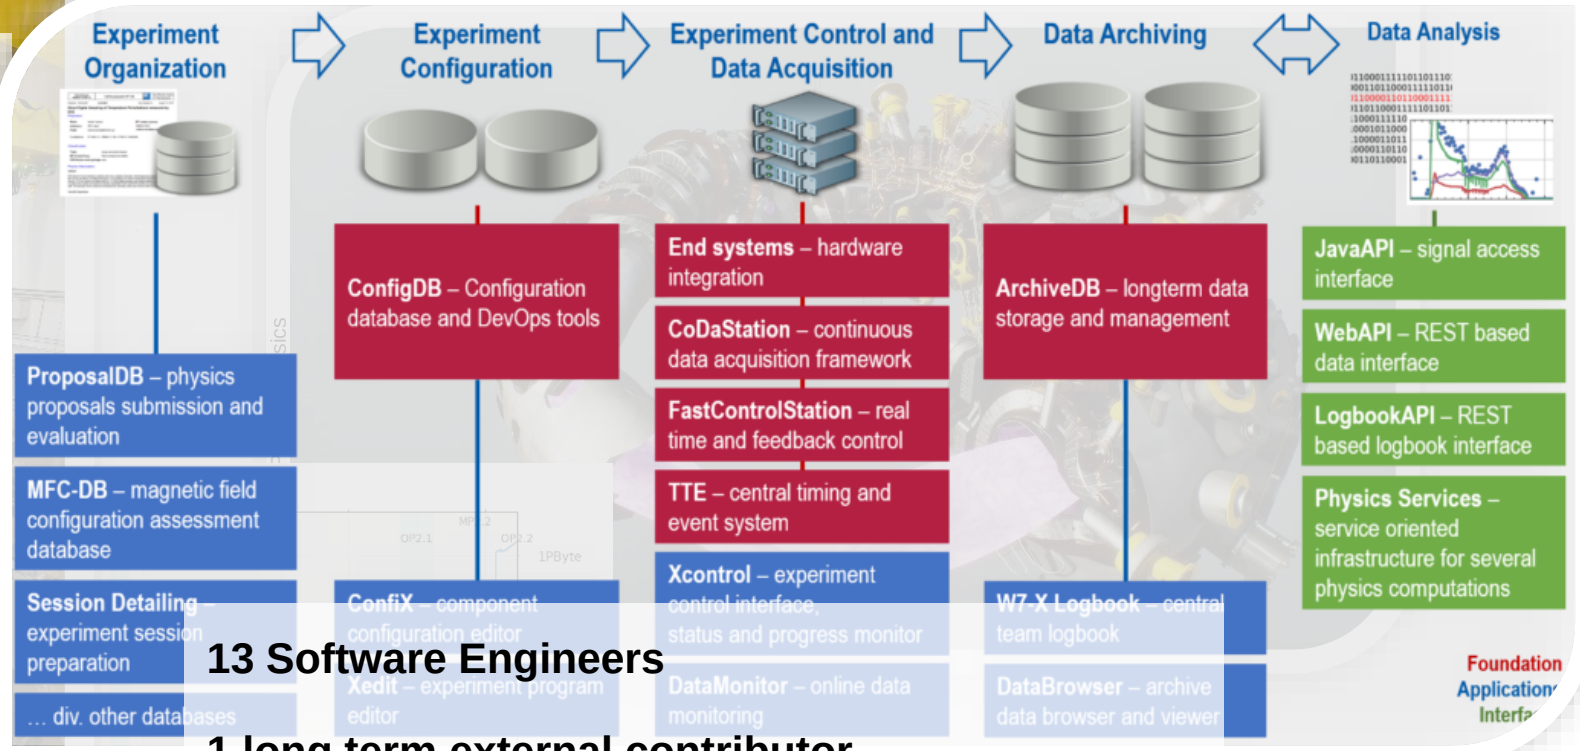

13 Software Engineers

1 long term external contributor

1 intern

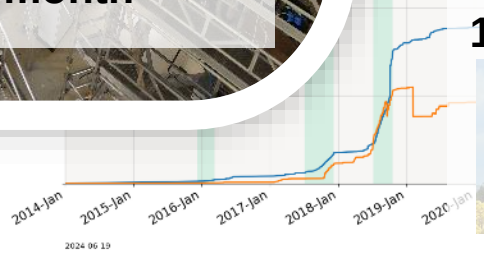

Agenda

Video compression

Digital images | Compression | Challenges | S.o.t. Art | FPV History | FPV Algorithm | FPV@W7X | Outlook | Standardization

Agenda

Video compression

- **Basics**
  - Digital images
  - Compression

Digital images | Compression | Challenges | S.o.t. Art | FPV History | FPV Algorithm | FPV@W7X | Outlook | Standardization

Agenda

Video compression

- **Basics**
  - Digital images
  - Compression
- **Challenges @W7-X**

Digital images | Compression | Challenges | S.o.t. Art | FPV History | FPV Algorithm | FPV@W7X | Outlook | Standardization

Agenda

Video compression

- **Basics**
  - Digital images
  - Compression
- **Challenges @W7-X**
- **State of the art**

Digital images | Compression | Challenges | S.o.t. Art | FPV History | FPV Algorithm | FPV@W7X | Outlook | Standardization

Agenda

Video compression

- **Basics**
  - Digital images
  - Compression
- **Challenges @W7-X**
- **State of the art**
- **Fusion Power Video**
  - History
  - Algorithm
  - Implementation @W7-X

Digital images | Compression | Challenges | S.o.t. Art | FPV History | FPV Algorithm | FPV@W7X | Outlook | Standardization

Agenda

Video compression

- **Basics**
  - Digital images
  - Compression
- **Challenges @W7-X**
- **State of the art**
- **Fusion Power Video**
  - History
  - Algorithm
  - Implementation @W7-X
- **Outlook**

Digital images | Compression | Challenges | S.o.t. Art | FPV History | FPV Algorithm | FPV@W7X | Outlook | Standardization

Agenda

Video compression

- **Basics**
  - Digital images
  - Compression
- **Challenges @W7-X**
- **State of the art**
- **Fusion Power Video**
  - History
  - Algorithm
  - Implementation @W7-X
- **Outlook**
- **Call for standardization**

Digital images | Compression | Challenges | S.o.t. Art | FPV History | FPV Algorithm | FPV@W7X | Outlook | Standardization

Basics – Digital images: arrays of relative intensity values

Basics – Digital images: arrays of relative intensity values

Basics – Digital images: arrays of relative intensity values

Basics – Digital images: arrays of relative intensity values

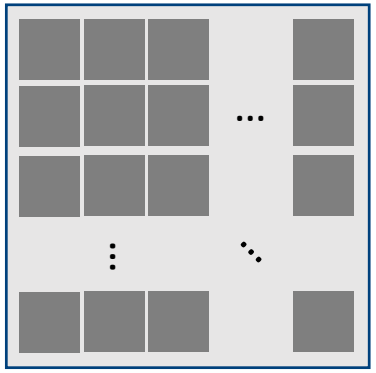

Basics – Digital images: arrays of relative intensity values

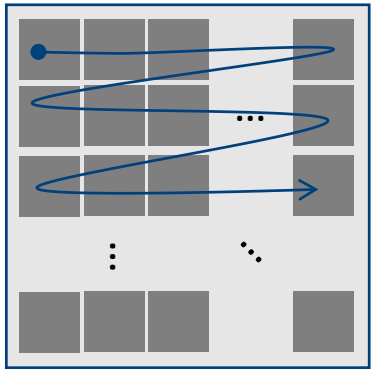

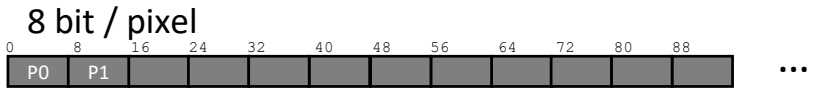

Basics – Digital images: arrays of relative intensity values

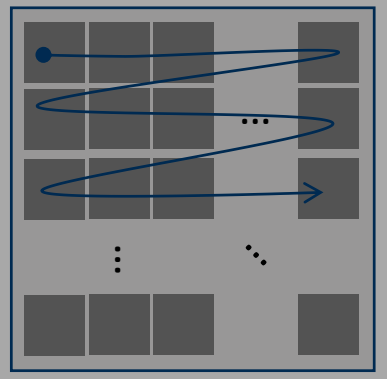

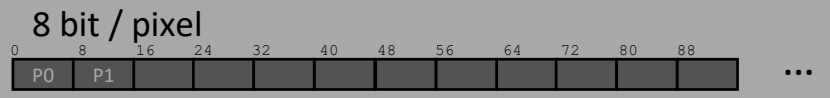

Its all relative

Basics – Digital images: arrays of relative intensity values

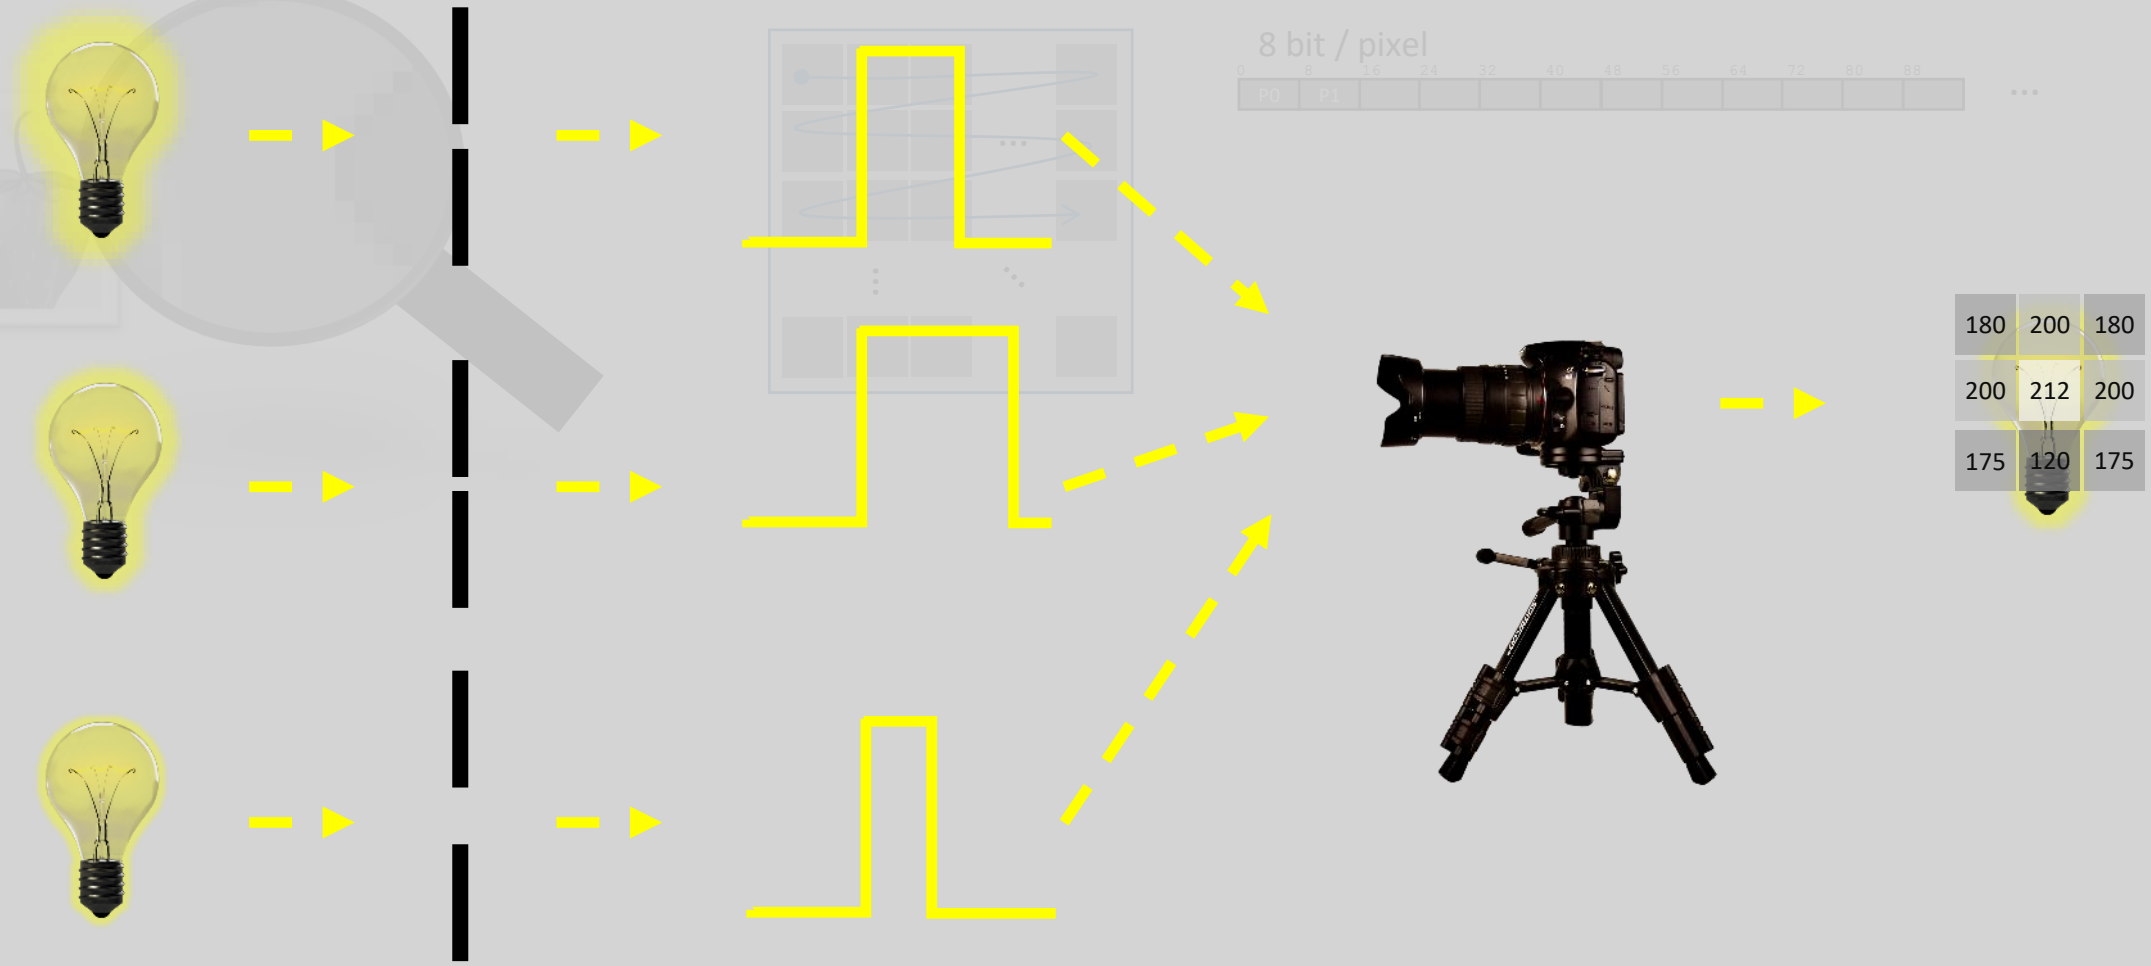

Its all relative

Basics – Digital images: arrays of relative intensity values

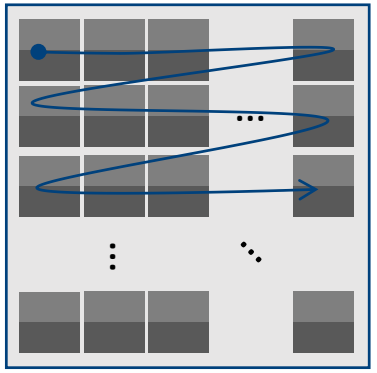

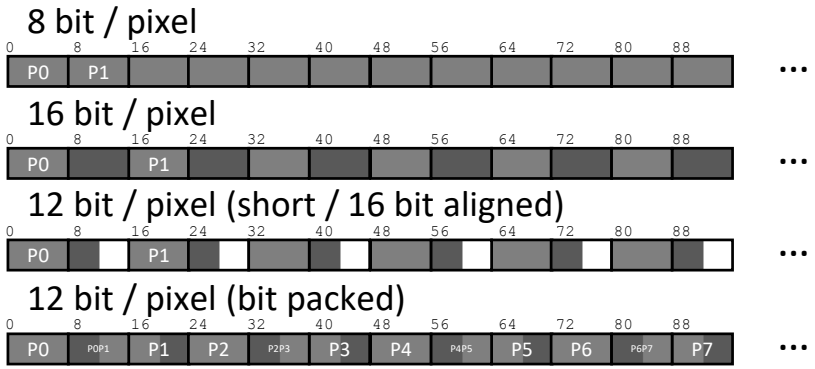

Basics – Digital images: arrays of relative intensity values

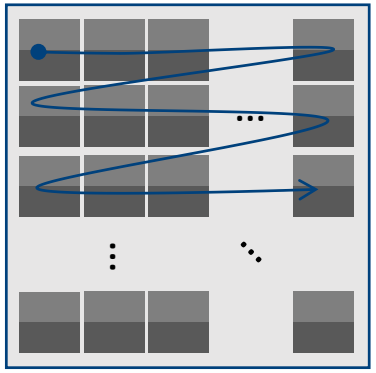

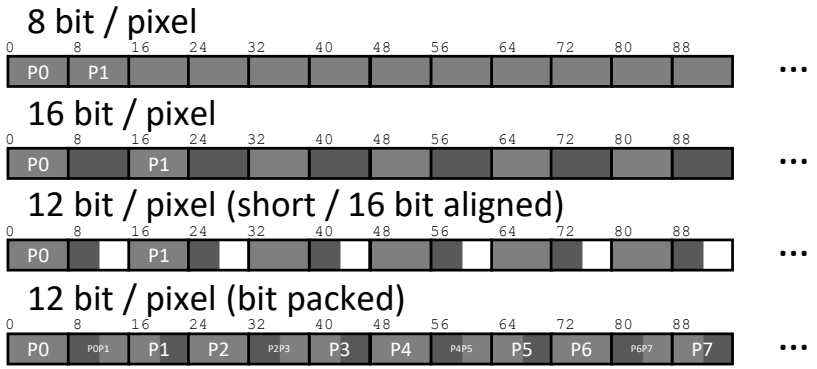

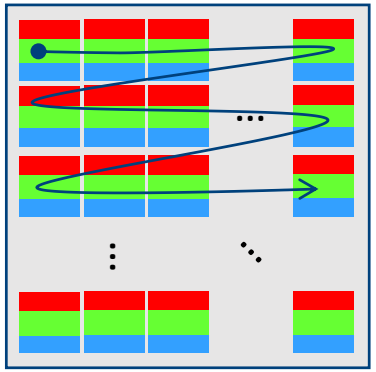

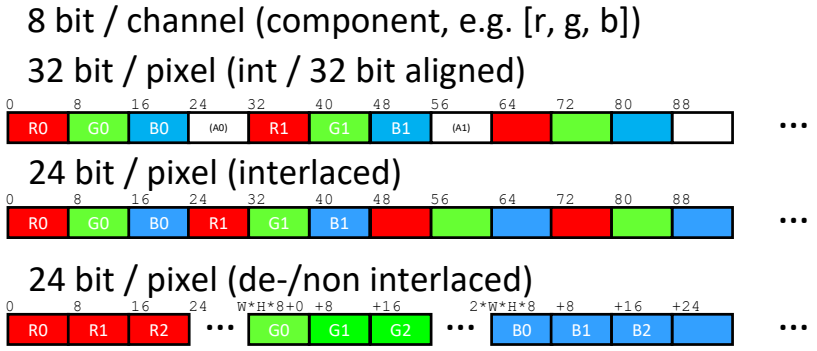

Basics – Digital images: arrays of relative intensity values

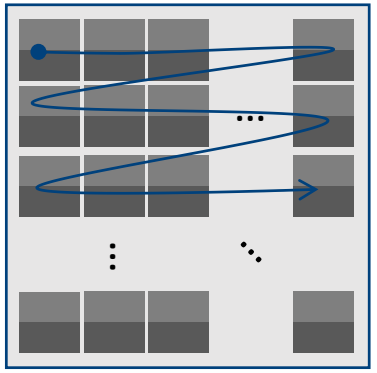

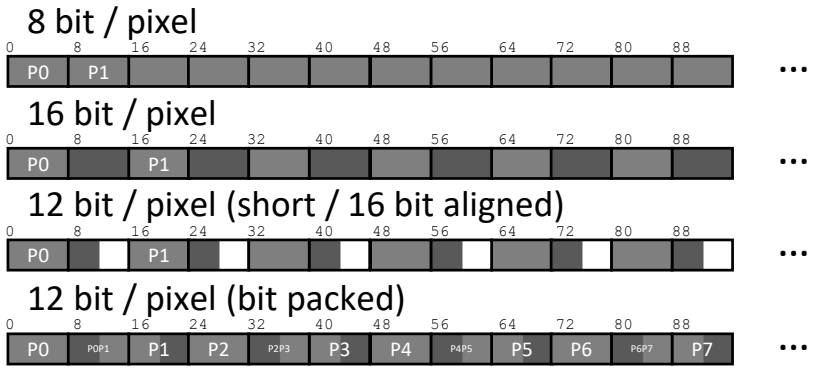

Basics – Compression schemes

Compression == transport the same information with less data

Basics – Compression schemes

Compression == transport the same information with less data

- Bit packing

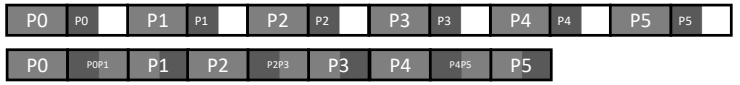

Basics – Compression schemes

Compression == transport the same information with less data

- Bit packing

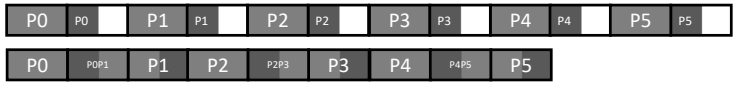

- Run length encoding

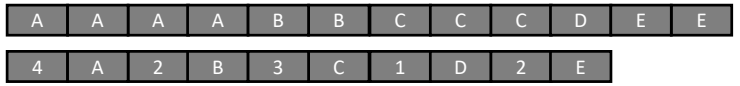

Digital images | **Compression** | Challenges | S.o.t. Art | FPV History | FPV Algorithm | FPV@W7X | Outlook | Standardization

Basics – Compression schemes

Compression == transport the same information with less data

- Bit packing

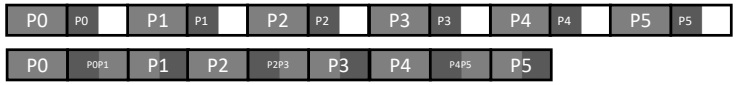

- Run length encoding

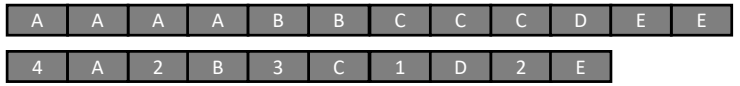

- Dictionary encoding

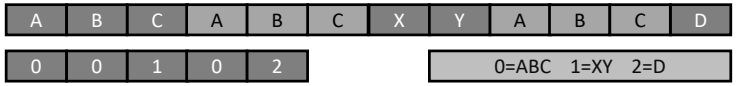

Basics – Compression schemes

Compression == transport the same information with less data

- Bit packing

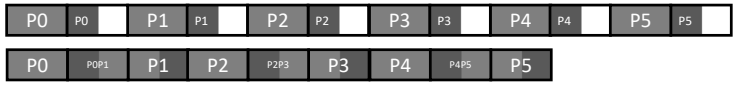

- Run length encoding

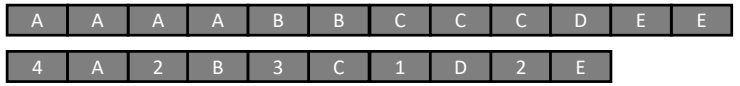

- Dictionary encoding

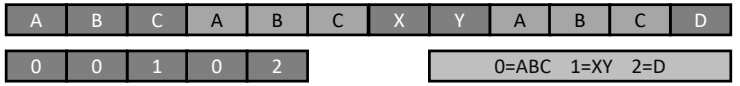

- Variable length encoding

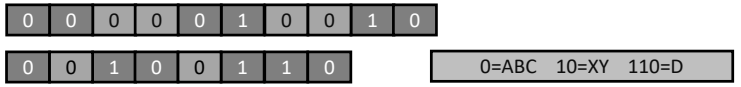

Basics – Compression schemes

Compression == transport the same information with less data

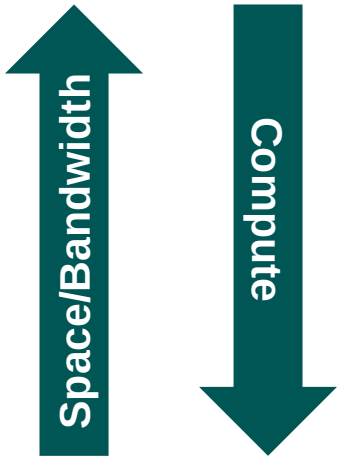

- Bit packing
- Run length encoding
- Dictionary encoding
- Variable length encoding

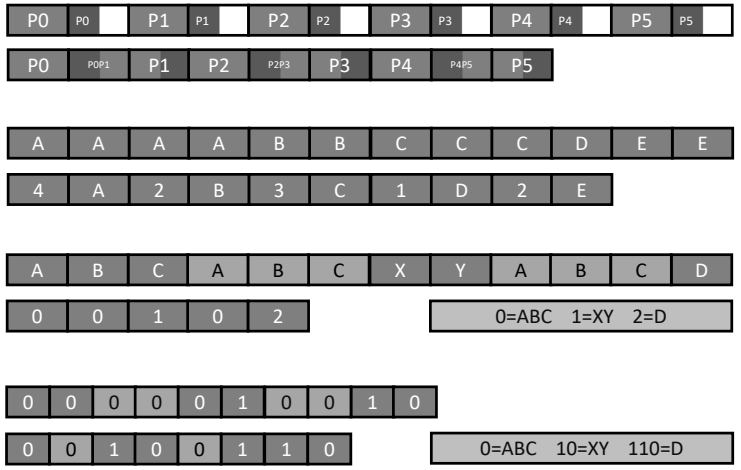

Basics – Compression schemes

Compression == transport the same information with less data

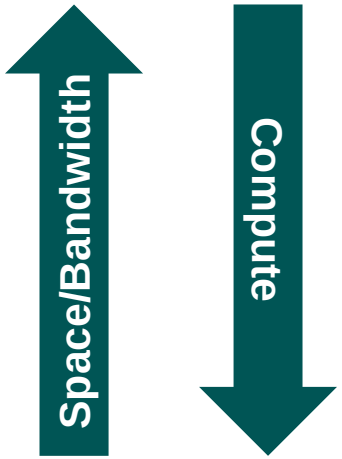

- Bit packing
- Run length encoding
- Dictionary encoding
- Variable length encoding

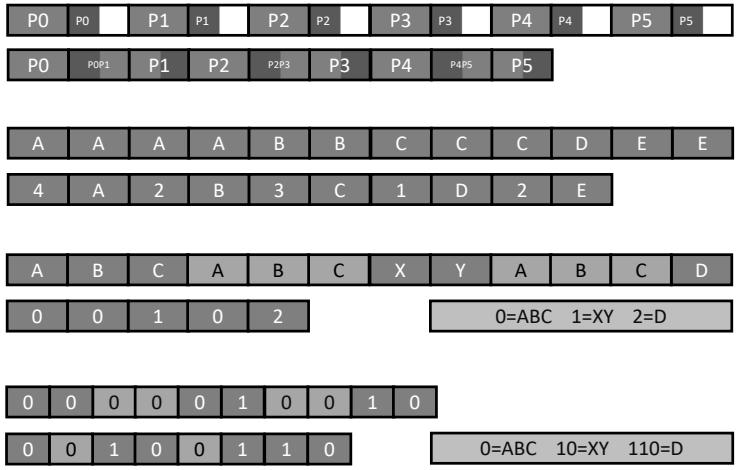

Reliance on patterns & their frequency

Basics – Compression schemes

Compression == transport the same information with less data

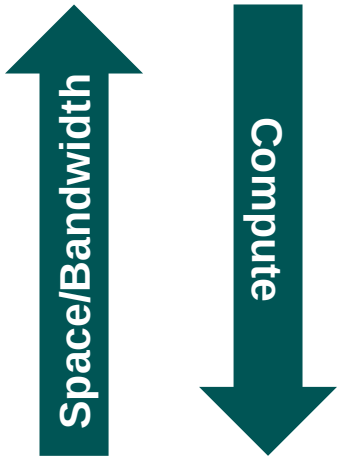

- Bit packing
- Run length encoding
- Dictionary encoding
- Variable length encoding

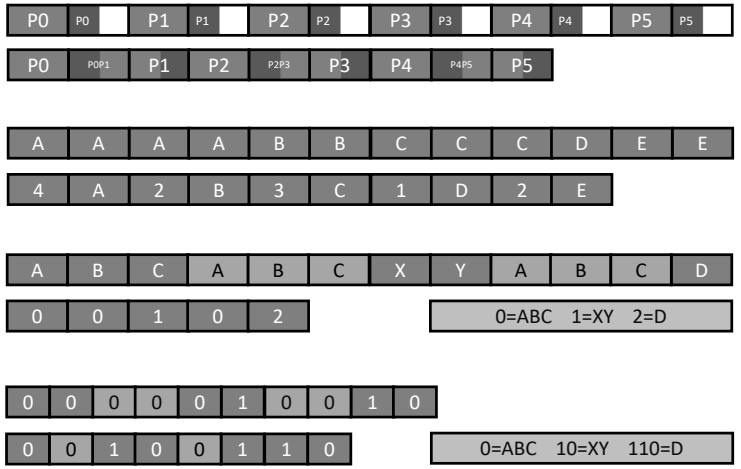

Reliance on patterns & their frequency

randomness/noise is the enemy

Basics – Compression schemes – loss & less

Lossless – visually lossless – lossy compression

Basics – Compression schemes – loss & less

Lossless – visually lossless – lossy compression

(lossy: the original cannot be identically reconstructed from the compressed)

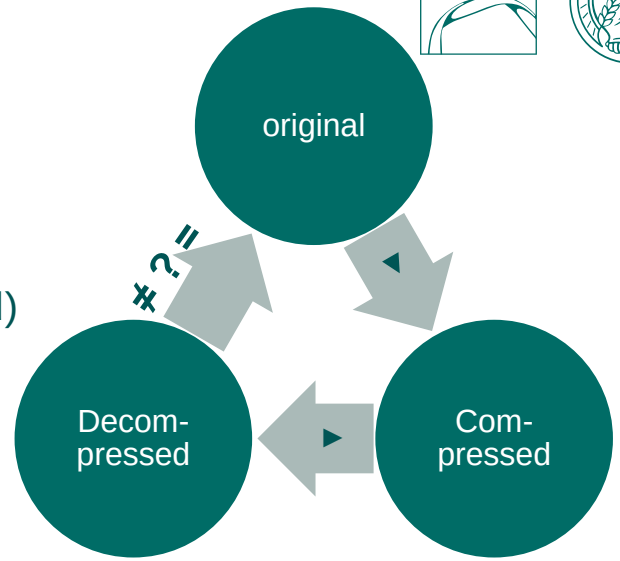

Digital images | **Compression** | Challenges | S.o.t. Art | FPV History | FPV Algorithm | FPV@W7X | Outlook | Standardization

Basics – Compression schemes – loss & less

Lossless – visually lossless – lossy compression

(lossy: the original cannot be identically reconstructed from the compressed)

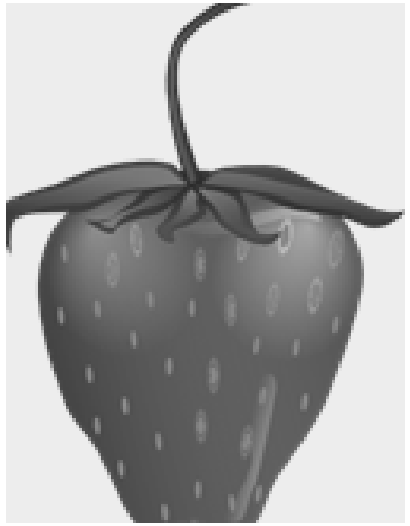

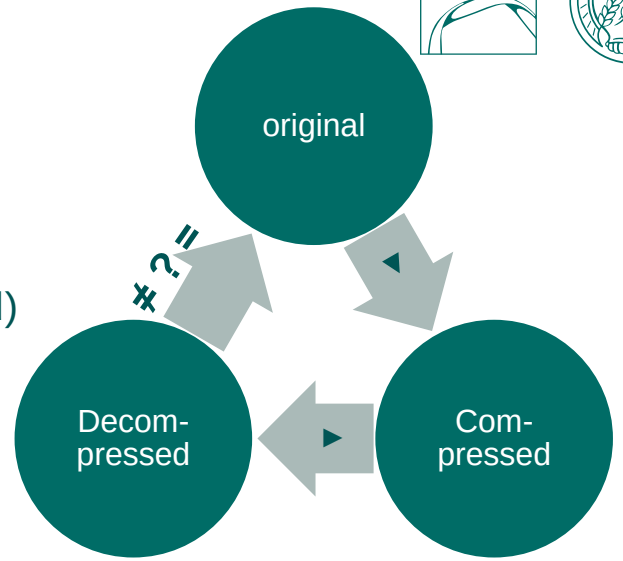

Basics – Compression schemes – loss & less

Lossless – visually lossless – lossy compression

(lossy: the original cannot be identically reconstructed from the compressed)

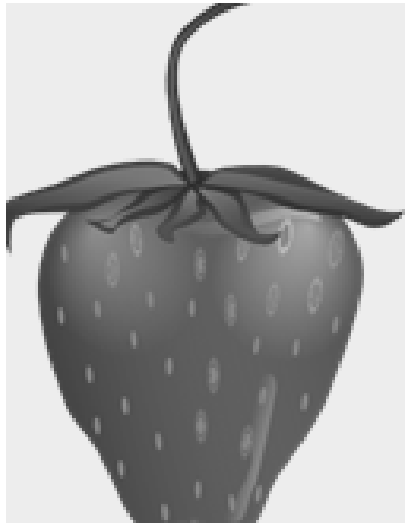

- Reduce bit depth

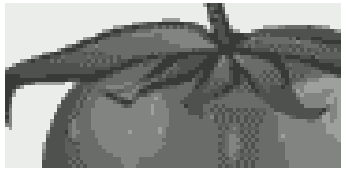

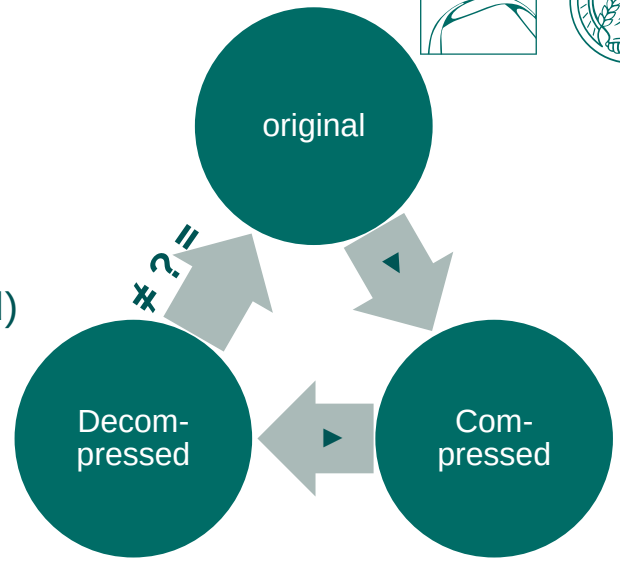

Basics – Compression schemes – loss & less

Lossless – visually lossless – lossy compression

(lossy: the original cannot be identically reconstructed from the compressed)

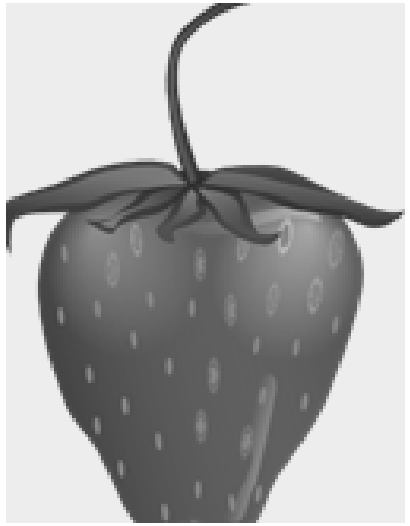

- Reduce bit depth
- Reduce resolution

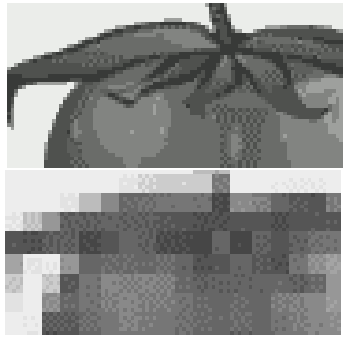

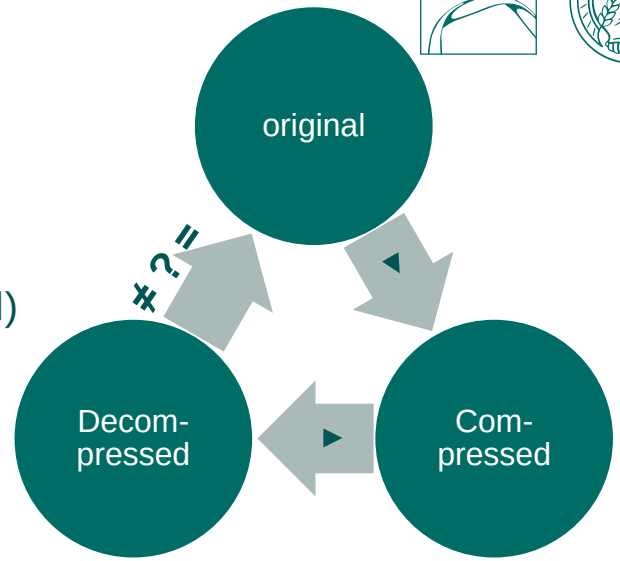

Basics – Compression schemes – loss & less

Lossless – visually lossless – lossy compression

(lossy: the original cannot be identically reconstructed from the compressed)

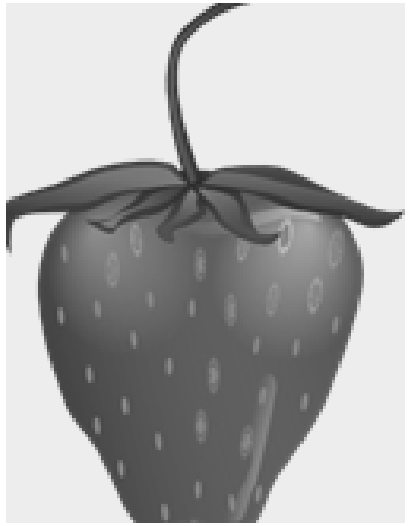

- Reduce bit depth
- Reduce resolution
- Reduce high frequencies

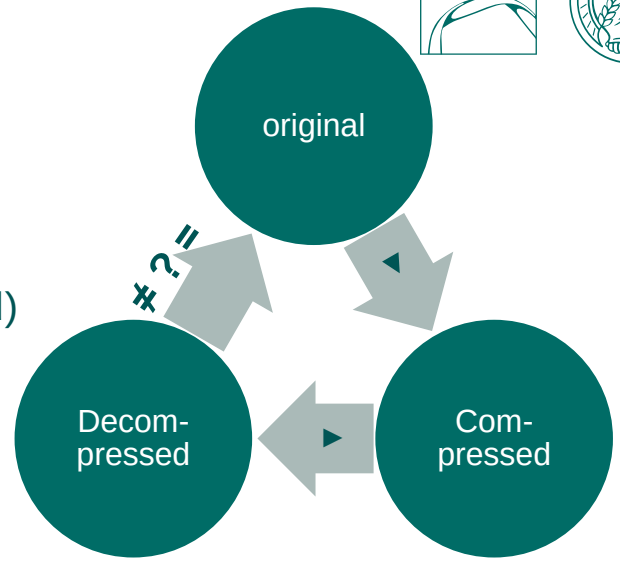

Basics – Compression schemes – loss & less

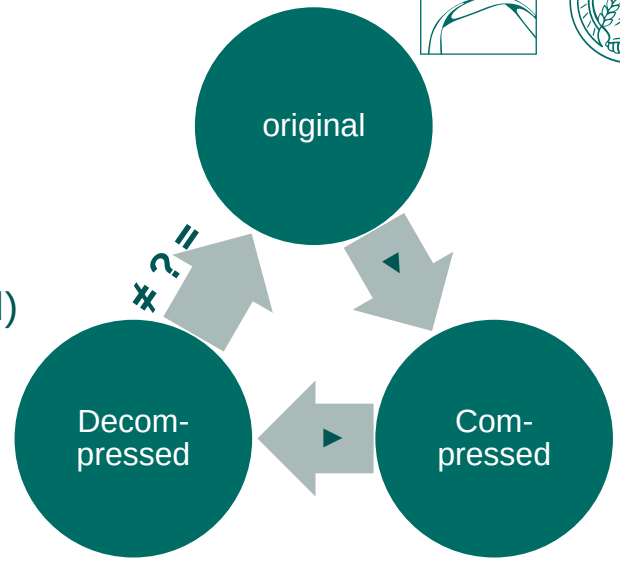

Lossless – visually lossless – lossy compression

(lossy: the original cannot be identically reconstructed from the compressed)

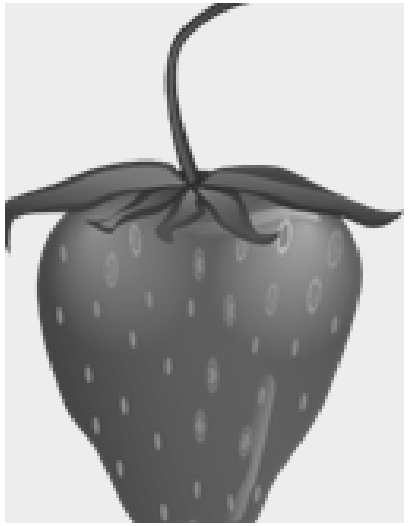

- Reduce bit depth
- Reduce resolution
- Reduce high frequencies

(visually lossless: apply the above, limited to a level not perceivable to the human eye)

Challenges at Wendelstein 7-X – camera level

White noise

Challenges at Wendelstein 7-X – camera level

White noise

- **Some cameras are very noisy**
- **White noise is *incompressible***

Challenges at Wendelstein 7-X – camera level

White noise

- **Some cameras are very noisy**
- **White noise is *incompressible***

Challenges at Wendelstein 7-X – camera level

White noise

- **Some cameras are very noisy**
- **White noise is *incompressible***

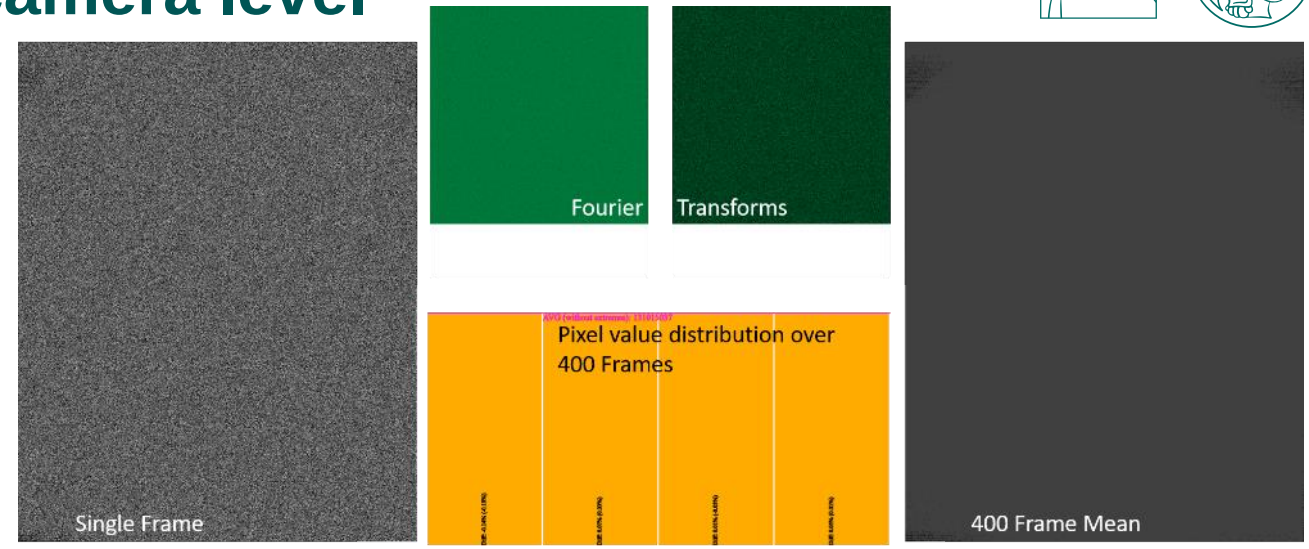

Challenges at Wendelstein 7-X – camera level

White noise

- **Some cameras are very noisy**
- **White noise is *incompressible***

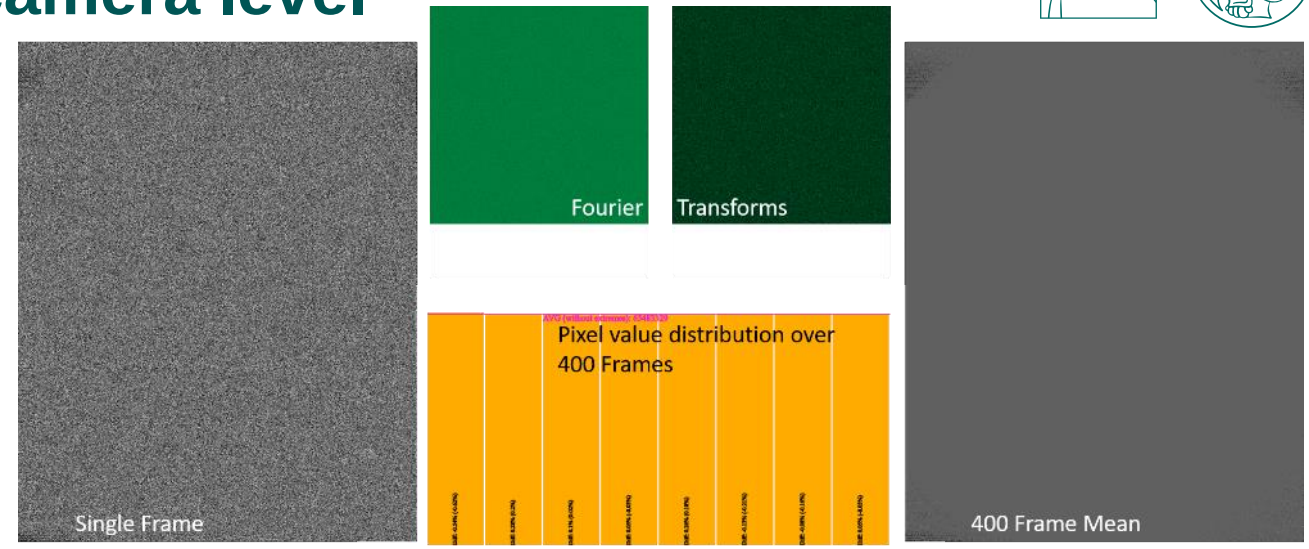

Challenges at Wendelstein 7-X – camera level

White noise

- **Some cameras are very noisy**
- **White noise is *incompressible***

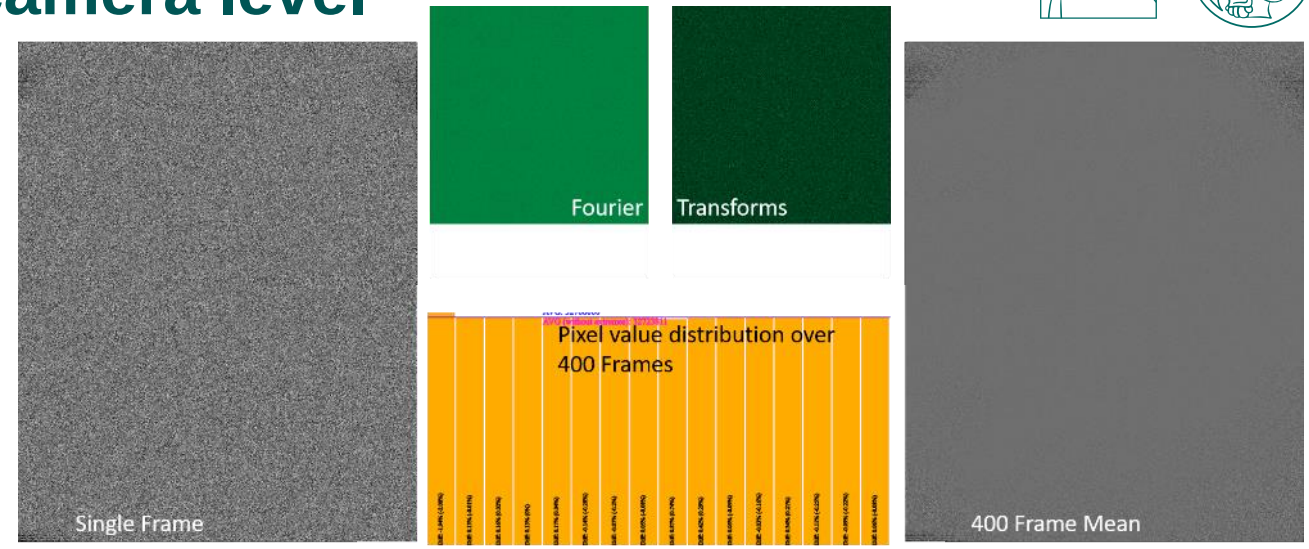

Challenges at Wendelstein 7-X – camera level

White noise

- **Some cameras are very noisy**
- **White noise is *incompressible***

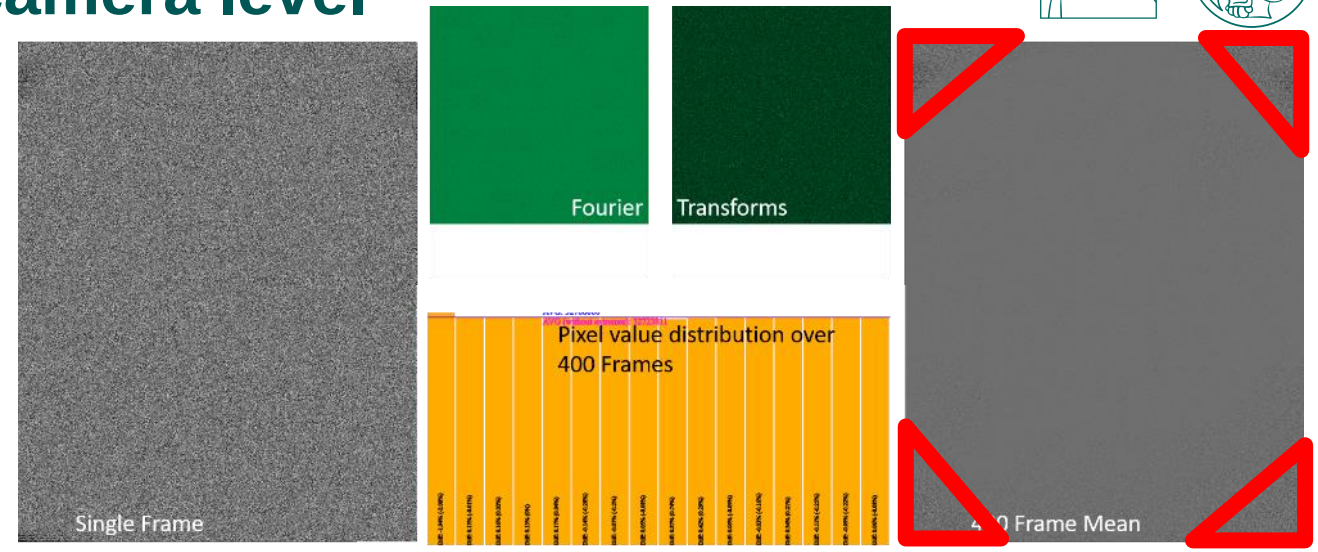

Challenges at Wendelstein 7-X – camera level

White noise

- Some cameras are very noisy
- White noise is ***incompressible***

(Non-)Alignment of pixel data

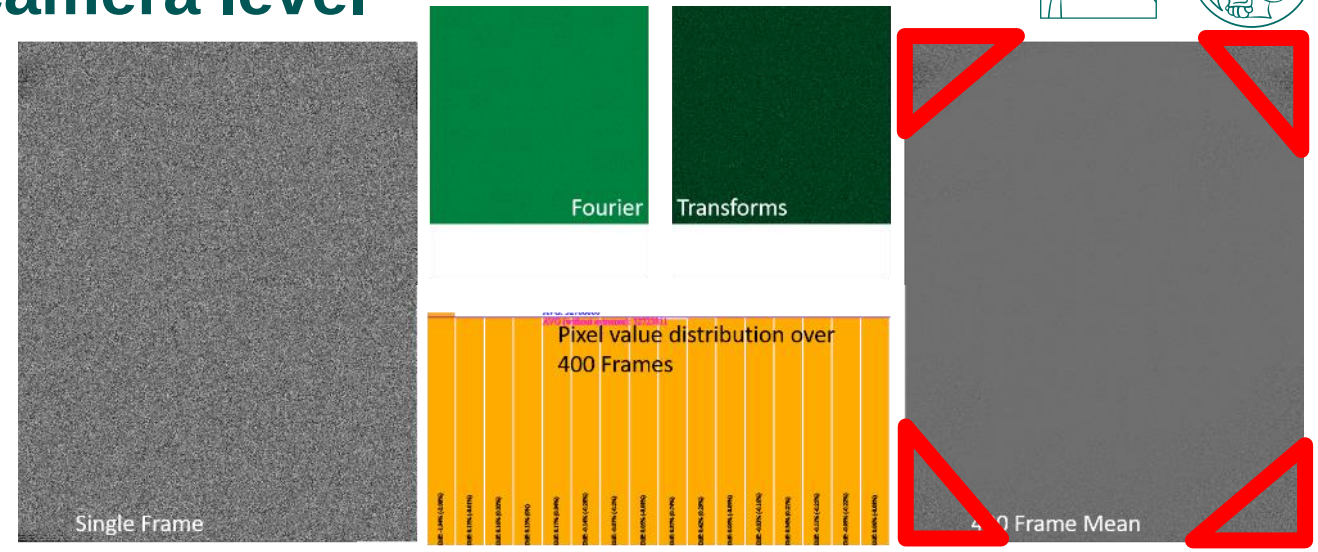

Challenges at Wendelstein 7-X – camera level

White noise

- Some cameras are very noisy
- White noise is ***incompressible***

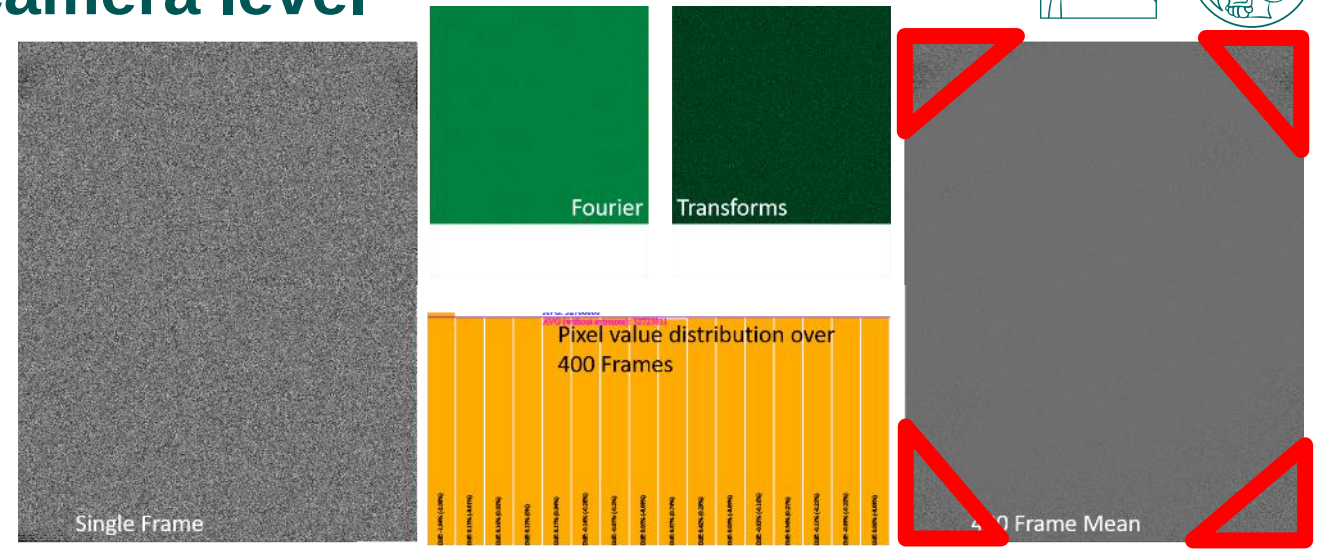

(Non-)Alignment of pixel data

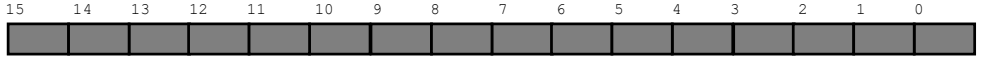

Challenges at Wendelstein 7-X – camera level

White noise

- Some cameras are very noisy
- White noise is ***incompressible***

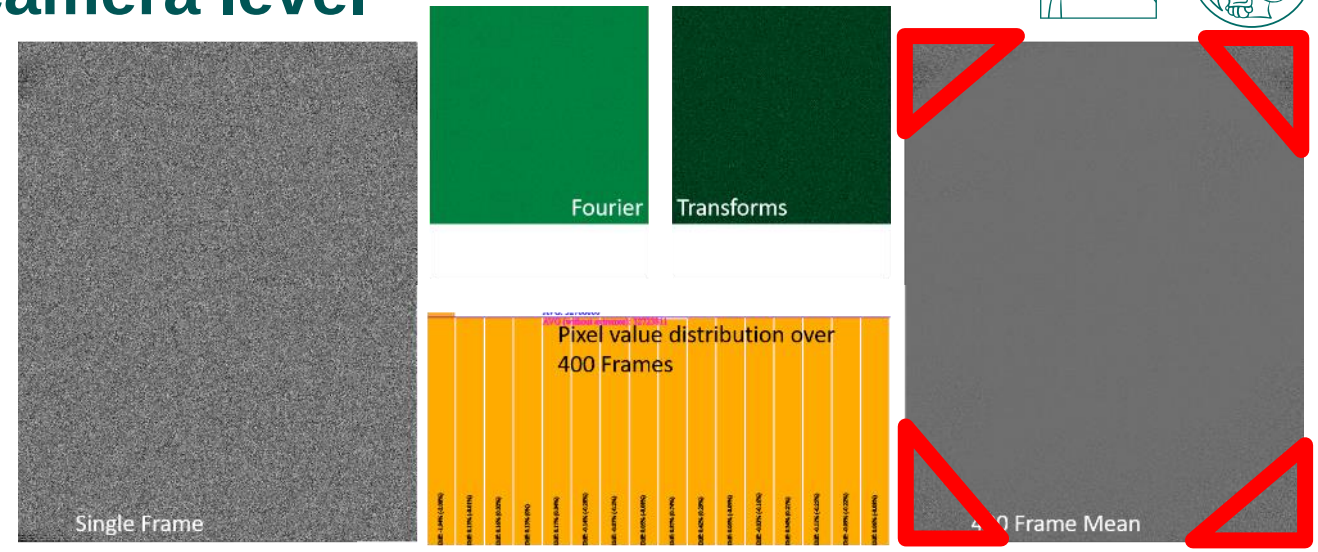

(Non-)Alignment of pixel data

Digital images | Compression | **Challenges** | S.o.t. Art | FPV History | FPV Algorithm | FPV@W7X | Outlook | Standardization

Challenges at Wendelstein 7-X – camera level

White noise

- Some cameras are very noisy
- White noise is ***incompressible***

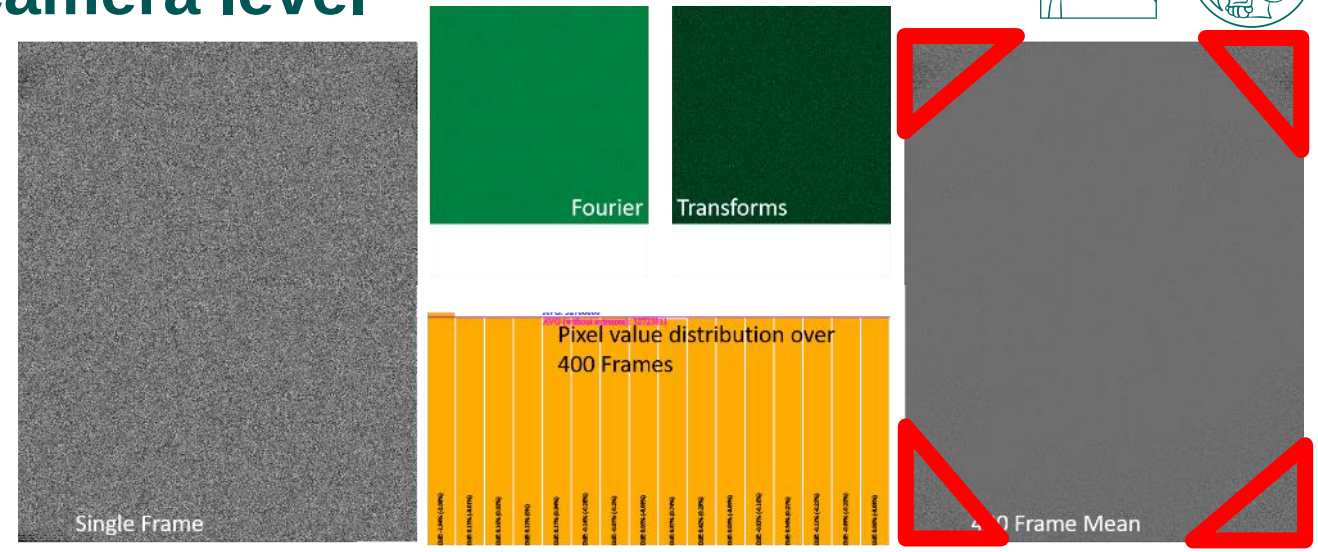

(Non-)Alignment of pixel data

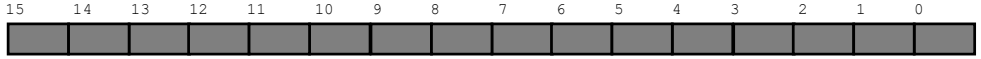

Digital images | Compression | **Challenges** | S.o.t. Art | FPV History | FPV Algorithm | FPV@W7X | Outlook | Standardization

Challenges at Wendelstein 7-X – camera level

White noise

- Some cameras are very noisy
- White noise is ***incompressible***

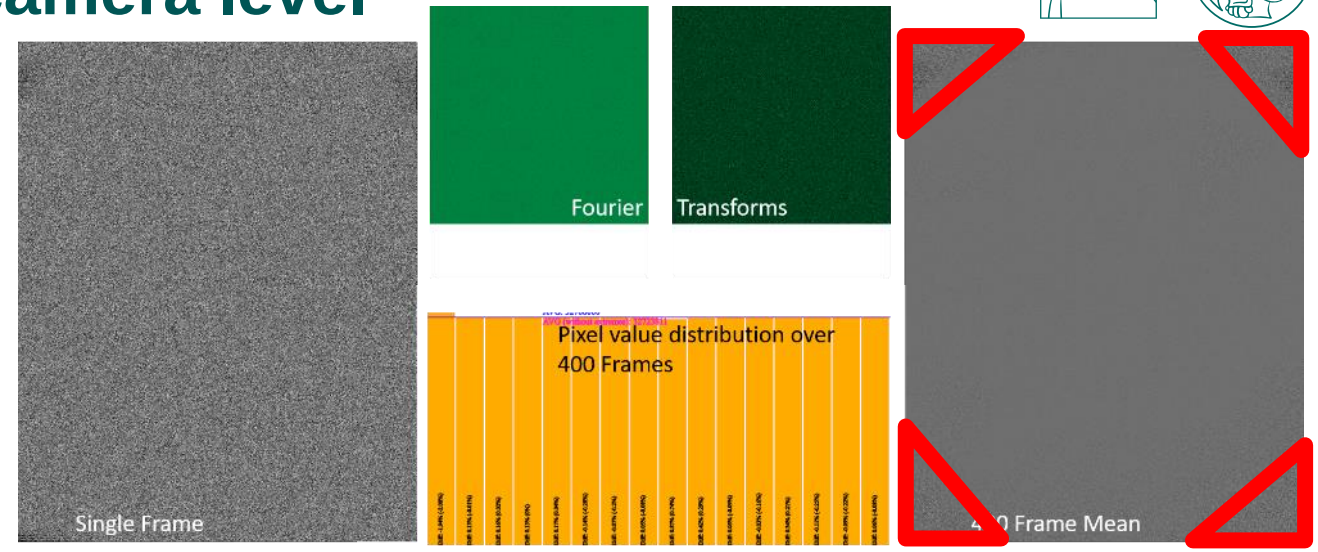

(Non-)Alignment of pixel data

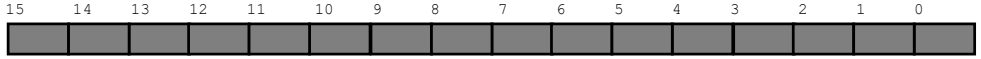

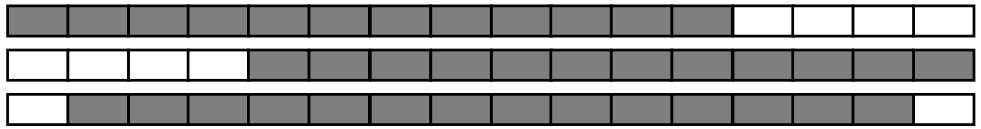

Digital images | Compression | **Challenges** | S.o.t. Art | FPV History | FPV Algorithm | FPV@W7X | Outlook | Standardization

Challenges at Wendelstein 7-X – camera level

White noise

- Some cameras are very noisy
- White noise is ***incompressible***

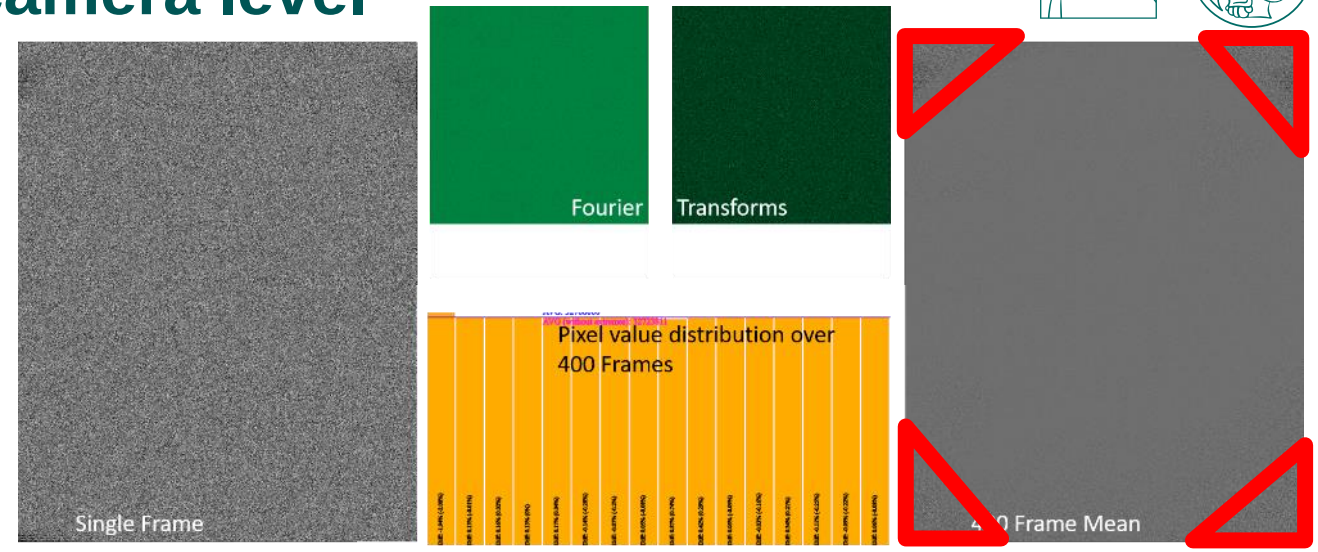

(Non-)Alignment of pixel data

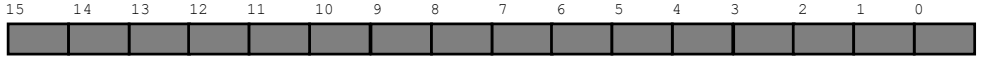

- You need to know

Challenges at Wendelstein 7-X – camera level

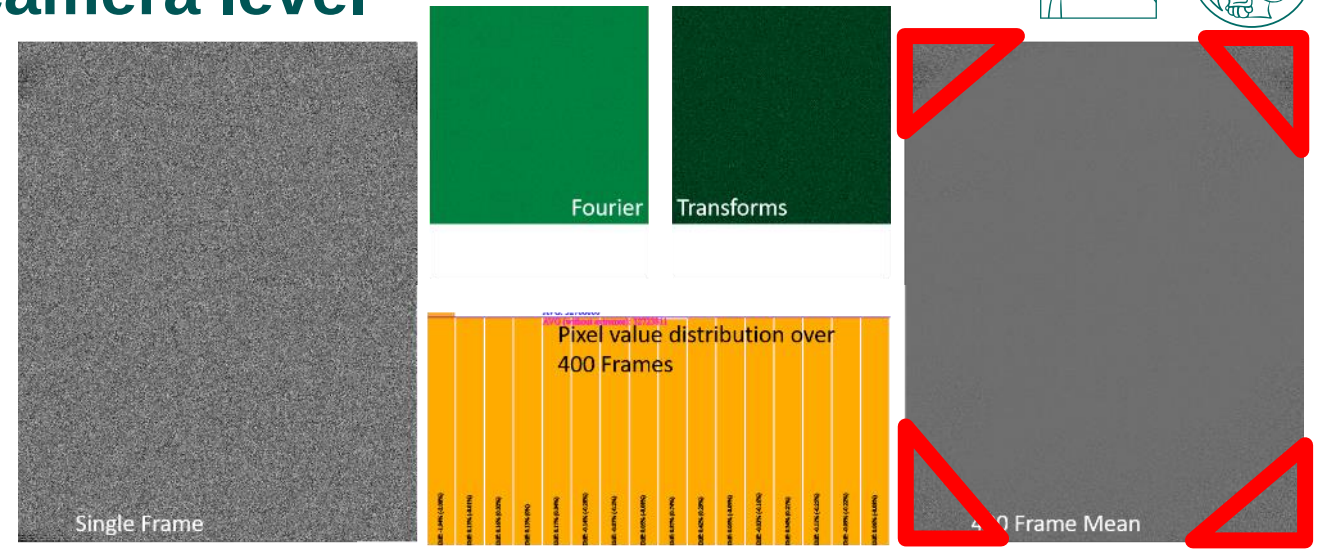

White noise

- Some cameras are very noisy
- White noise is ***incompressible***

(Non-)Alignment of pixel data

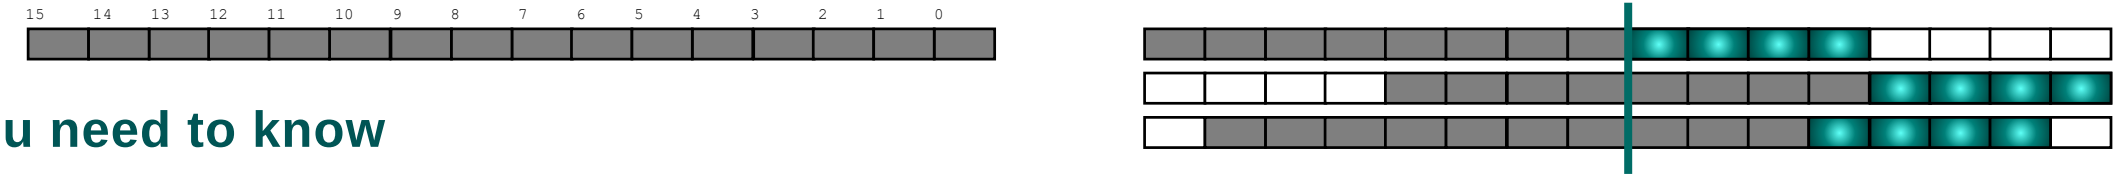

- You need to know
- Complicates finding the non-noisy part & assess the relativity

Challenges at Wendelstein 7-X – camera level

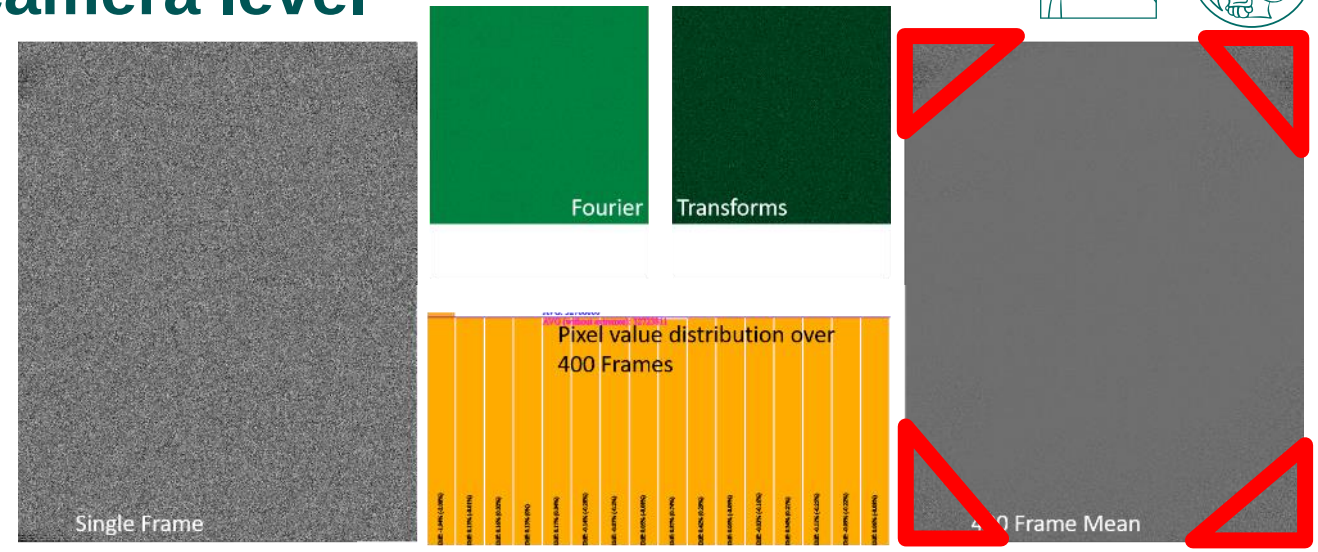

White noise

- Some cameras are very noisy
- White noise is ***incompressible***

(Non-)Alignment of pixel data

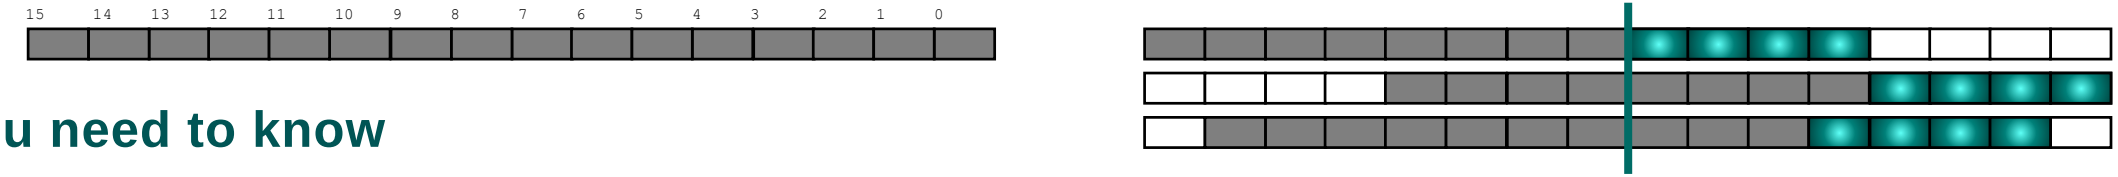

- You need to know
- Complicates finding the non-noisy part & assess the relativity

Requirement for lossless

Challenges at Wendelstein 7-X – camera level

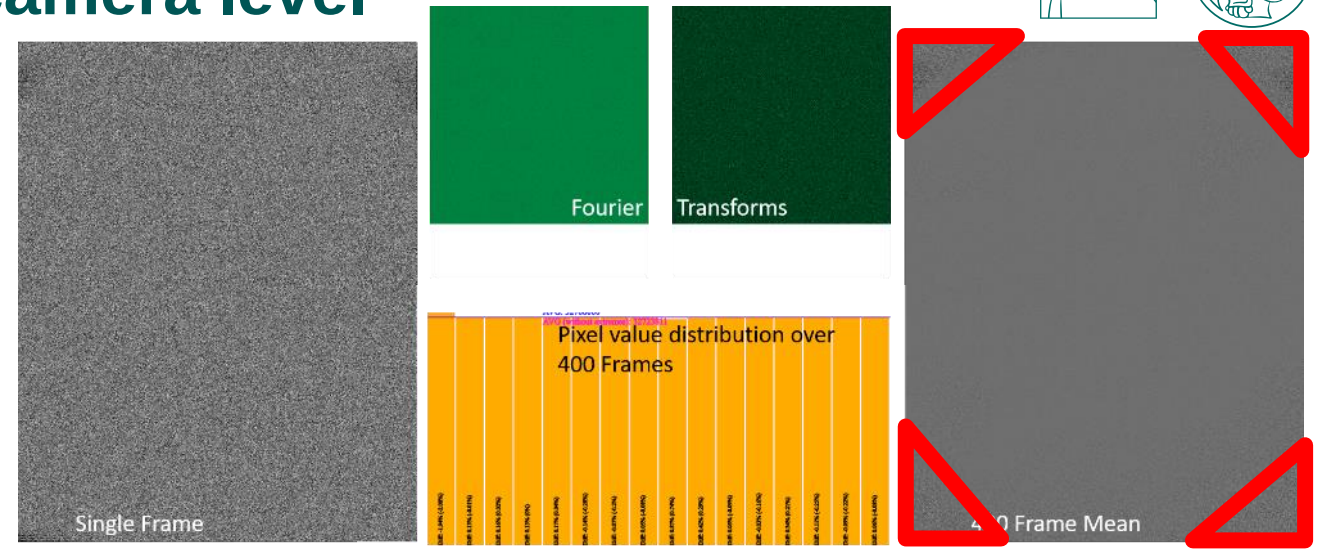

White noise

- Some cameras are very noisy
- White noise is ***incompressible***

(Non-)Alignment of pixel data

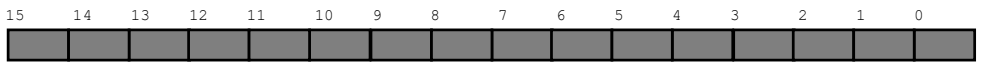

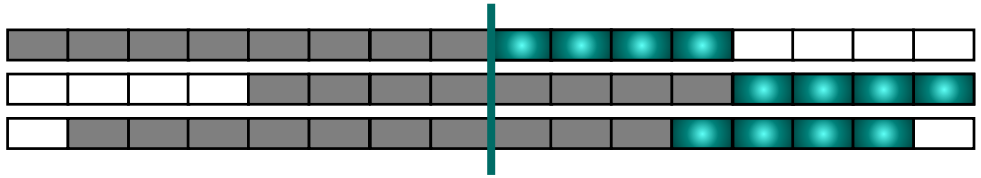

- You need to know
- Complicates finding the non-noisy part & assess the relativity

Requirement for lossless

- Lossy compression not only looses scientific data, ***but changes it***

Challenges at Wendelstein 7-X – camera level

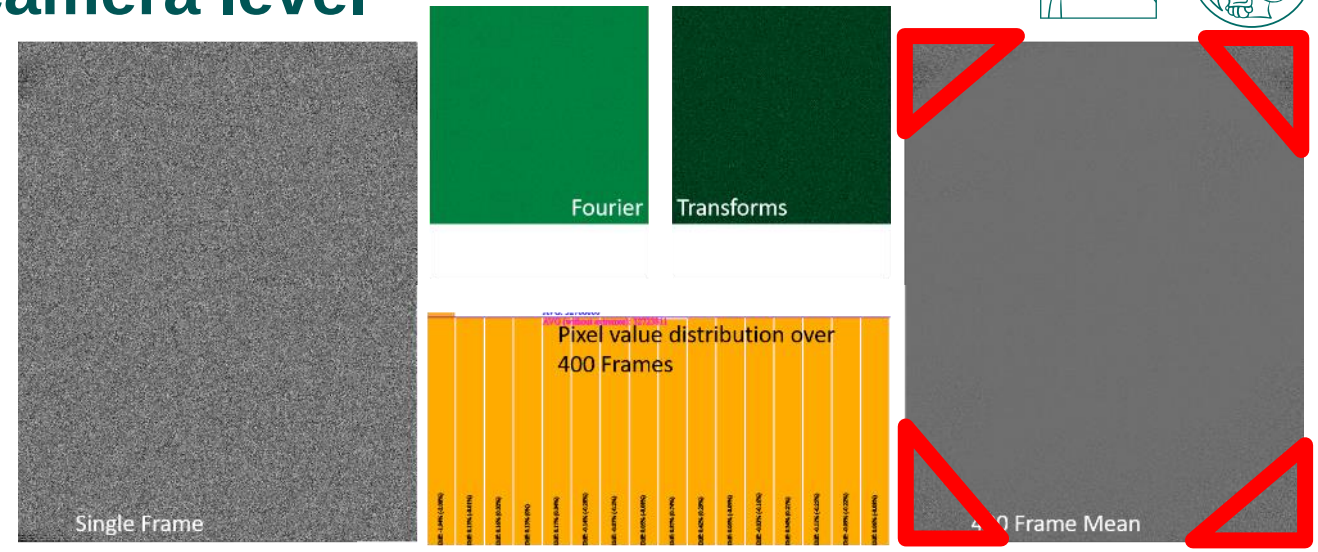

White noise

- Some cameras are very noisy
- White noise is ***incompressible***

(Non-)Alignment of pixel data

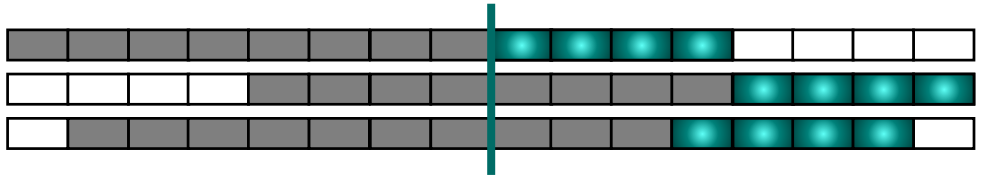

- You need to know
- Complicates finding the non-noisy part & assess the relativity

Requirement for lossless

- Lossy compression not only looses scientific data, ***but changes it***
- ***There also could be a Nobel Prize hidden in those bits we lose…***

Challenges at Wendelstein 7-X – machine level

Long Pulses & lots of Cameras

Challenges at Wendelstein 7-X – machine level

Long Pulses & lots of Cameras

- **Stellarators (in theory) support continuous operation**
  - **Goal of W7-X: 30 minutes**

Challenges at Wendelstein 7-X – machine level

Long Pulses & lots of Cameras

- **Stellarators (in theory) support continuous operation**
  - **Goal of W7-X: 30 minutes**

➢ **Archiving/compression required in (near) real time**

Challenges at Wendelstein 7-X – machine level

Long Pulses & lots of Cameras

- **Stellarators (in theory) support continuous operation**
  - **Goal of W7-X: 30 minutes**

➢ **Archiving/compression required in (near) real time**

- **Present: ~100 cameras – .8 to 5.5 MP – 25 to 400 FPS (+ analyzed data)**

Challenges at Wendelstein 7-X – machine level

Long Pulses & lots of Cameras

- **Stellarators (in theory) support continuous operation**
  - **Goal of W7-X: 30 minutes**

➢ **Archiving/compression required in (near) real time**

- **Present: ~100 cameras – .8 to 5.5 MP – 25 to 400 FPS (+ analyzed data)**
- **Example OP2.1: Divertor thermography + Halpha + EDICAM:**

Challenges at Wendelstein 7-X – machine level

Long Pulses & lots of Cameras

- Stellarators (in theory) support continuous operation
  - Goal of W7-X: 30 minutes

➢ Archiving/compression required in (near) real time

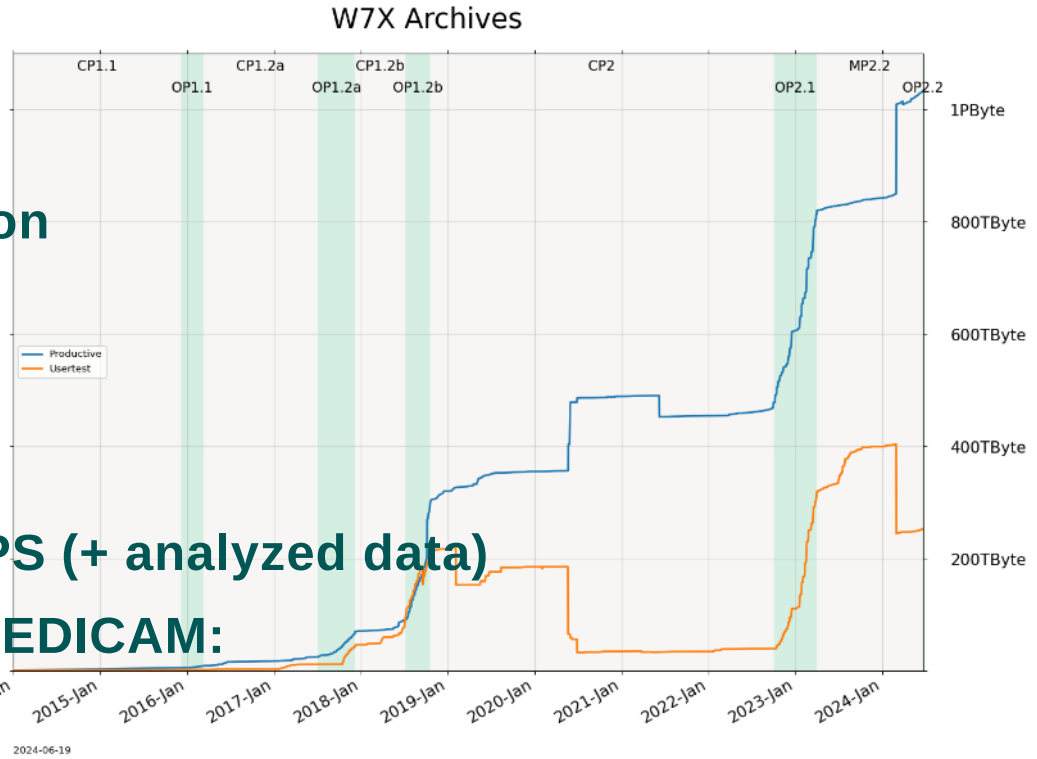

- Present: ~100 cameras – .8 to 5.5 MP – 25 to 400 FPS (+ analyzed data)
- Example OP2.1: Divertor thermography + Halpha + EDICAM:

35 cameras + 12x thermal: ~.8PByte raw data

Challenges at Wendelstein 7-X – machine level

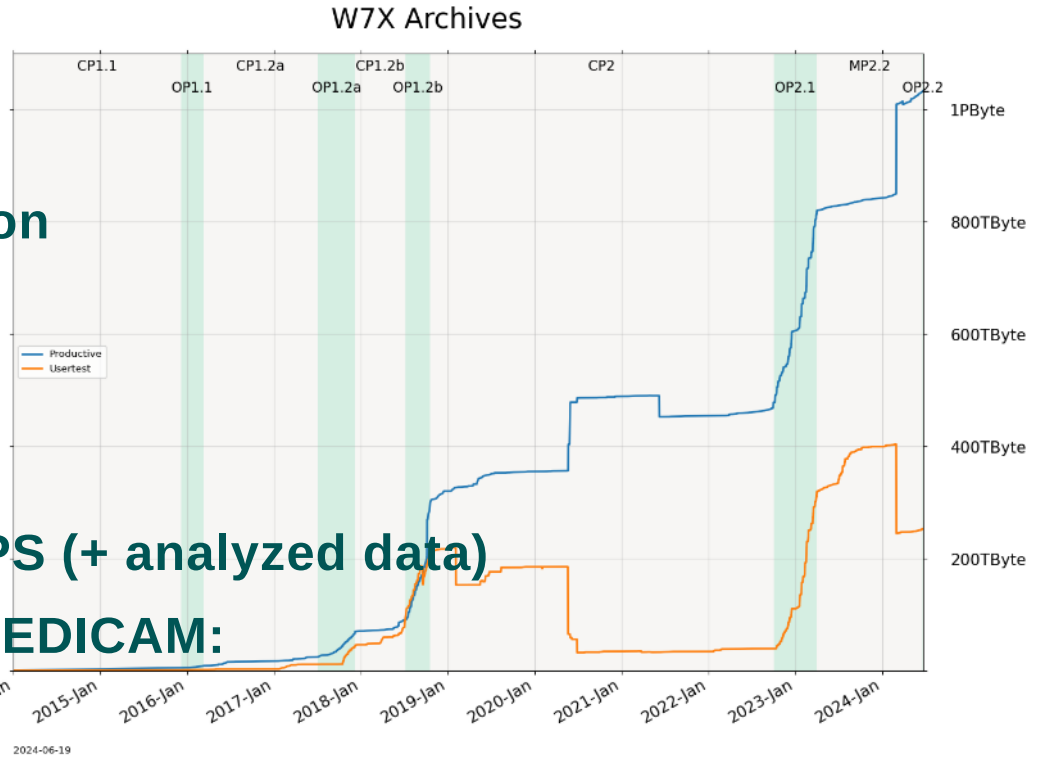

Long Pulses & lots of Cameras

- Stellarators (in theory) support continuous operation
  - Goal of W7-X: 30 minutes

➢ Archiving/compression required in (near) real time

- Present: ~100 cameras – .8 to 5.5 MP – 25 to 400 FPS (+ analyzed data)
- Example OP2.1: Divertor thermography + Halpha + EDICAM:

35 cameras + 12x thermal: ~.8PByte raw data

Data swamp

- Risk: relevant potions are hard to find -> a lot of archived data is rarely/never read

Challenges at Wendelstein 7-X – machine level

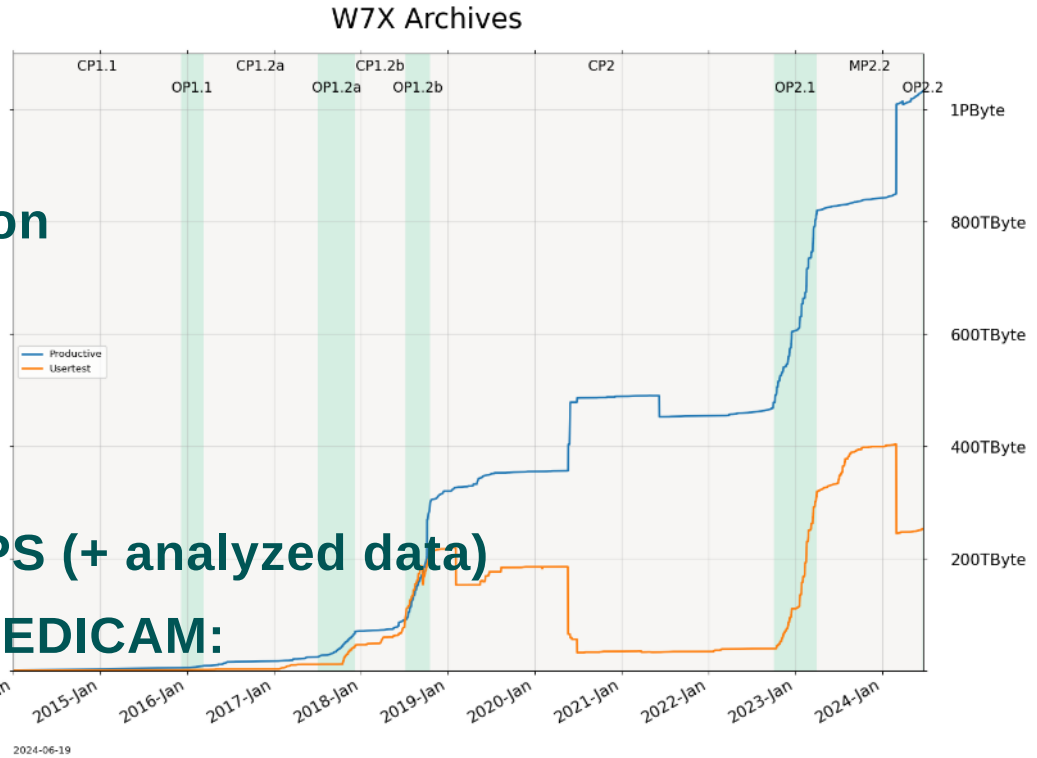

Long Pulses & lots of Cameras

- Stellarators (in theory) support continuous operation
  - Goal of W7-X: 30 minutes

➢ Archiving/compression required in (near) real time

- Present: ~100 cameras – .8 to 5.5 MP – 25 to 400 FPS (+ analyzed data)
- Example OP2.1: Divertor thermography + Halpha + EDICAM:

35 cameras + 12x thermal: ~.8PByte raw data

Data swamp

- Risk: relevant potions are hard to find -> a lot of archived data is rarely/never read

➢ Fast (overview) access needed to support filtering

State of the art

How do people archive images / videos for scientific use?

State of the art

How do people archive images / videos for scientific use?

- **Store them uncompressed (aren’t that many & storage is cheap)**

State of the art

How do people archive images / videos for scientific use?

- **Store them uncompressed (aren’t that many & storage is cheap)**
- **They don’t (archive them)**

State of the art

How do people archive images / videos for scientific use?

- **Store them uncompressed (aren't that many & storage is cheap)**
- **They don't (archive them)**
- **Throw most away (after selecting an important few)**

State of the art

How do people archive images / videos for scientific use?

- **Store them uncompressed (aren't that many & storage is cheap)**
- **They don't (archive them)**
- **Throw most away (after selecting an important few)**
- **Compress utilizing general bit stream compressors (or compressing FS)**

State of the art

How do people archive images / videos for scientific use?

- **Store them uncompressed (aren't that many & storage is cheap)**
- **They don't (archive them)**
- **Throw most away (after selecting an important few)**
- **Compress utilizing general bit stream compressors (or compressing FS)**
- **Store lossy/visually lossless, utilizing (HW supported) standard video compressors**

State of the art

How do people archive images / videos for scientific use?

- **Store them uncompressed (aren't that many & storage is cheap)**
- **They don't (archive them)**
- **Throw most away (after selecting an important few)**
- **Compress utilizing general bit stream compressors (or compressing FS)**
- **Store lossy/visually lossless, utilizing (HW supported) standard video compressors**
- **Get creative with lossless H264**

State of the art

How do people archive images / videos for scientific use?

- **Store them uncompressed (aren't that many & storage is cheap)**
- **They don't (archive them)**
- **Throw most away (after selecting an important few)**
- **Compress utilizing general bit stream compressors (or compressing FS)**
- **Store lossy/visually lossless, utilizing (HW supported) standard video compressors**
- **Get creative with lossless H264**
  - **Cut bit depths down to 8-10 bpp grayscale**

State of the art

How do people archive images / videos for scientific use?

- **Store them uncompressed (aren't that many & storage is cheap)**
- **They don't (archive them)**
- **Throw most away (after selecting an important few)**
- **Compress utilizing general bit stream compressors (or compressing FS)**
- **Store lossy/visually lossless, utilizing (HW supported) standard video compressors**
- **Get creative with lossless H264**
  - **Cut bit depths down to 8-10 bpp grayscale**
  - **(Mis-)use chroma planes of YUV**

State of the art

How do people archive images / videos for scientific use?

- **Store them uncompressed (aren't that many & storage is cheap)**
- **They don't (archive them)**
- **Throw most away (after selecting an important few)**
- **Compress utilizing general bit stream compressors (or compressing FS)**
- **Store lossy/visually lossless, utilizing (HW supported) standard video compressors**
- **Get creative with lossless H264**
  - **Cut bit depths down to 8-10 bpp grayscale**
  - **(Mis-)use chroma planes of YUV**

And you?

State of the art

What do we do at W7-X?

State of the art

What do we do at W7-X?

- **Normalize to intensity maximum at 65535**

State of the art

What do we do at W7-X?

- **Normalize to intensity maximum at 65535**
- **Compress losslessly, in real time, without specific hardware acceleration**

State of the art

What do we do at W7-X?

- **Normalize to intensity maximum at 65535**
- **Compress losslessly, in real time, without specific hardware acceleration**
- **Optimize storage format for fast read out of most important/valuable data**

State of the art

What do we do at W7-X?

- **Normalize to intensity maximum at 65535**
- **Compress losslessly, in real time, without specific hardware acceleration**
- **Optimize storage format for fast read out of most important/valuable data**
- **… and even faster preview**

State of the art

What do we do at W7-X?

- Normalize to intensity maximum at 65535
- Compress losslessly, in real time, without specific hardware acceleration
- Optimize storage format for fast read out of most important/valuable data
- … and even faster preview

Compression Ratio: 36% (- 58%*)
Savings: 0.5 Petabytes of archive storage

+ no network upgrade necessary

***58% due to configuration error: data not normalized**

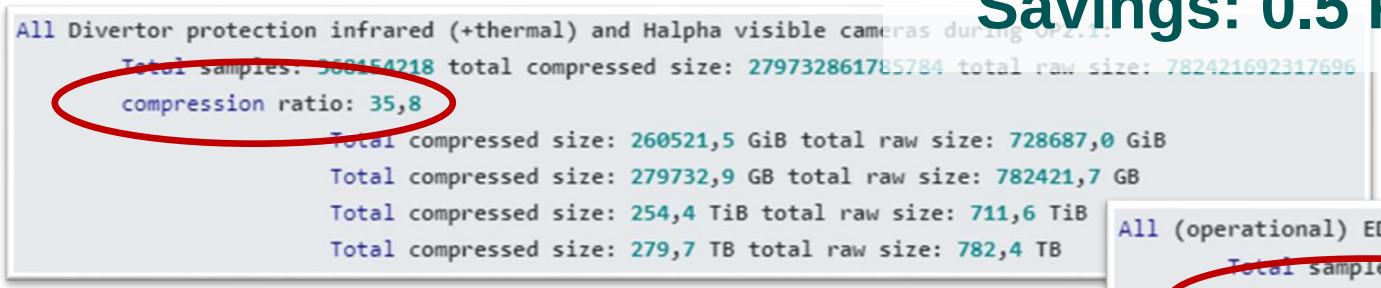

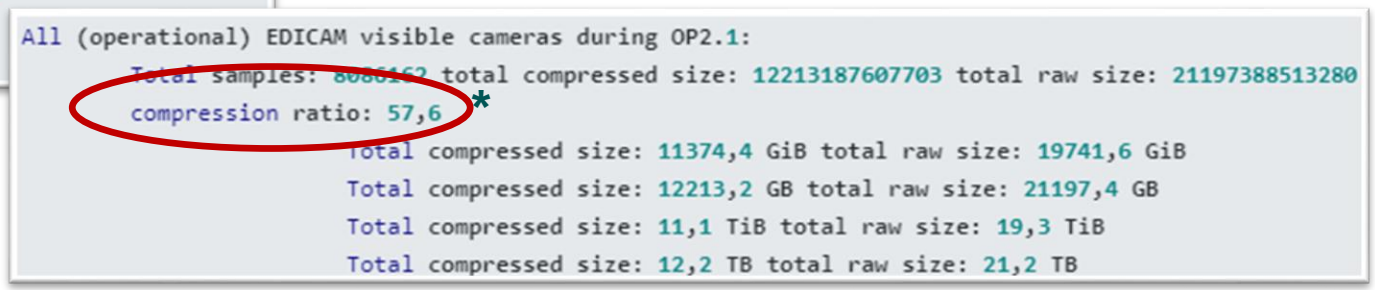

Fusion Power Video – A short history (I)

We need to compress losslessly – what options do we have?

Fusion Power Video – A short history (I)

We need to compress losslessly – what options do we have?

- **Bit stream compressors: ZSTD on fast – 60% of real time (RT), 60% compression**

Reference: Intel Xeon, 3.4GHz, 24 cores; 128GB RAM, Geforce GTX 2060 6GB
X% compression ⇔ compressed size is X% of raw size
data from 2020

Fusion Power Video – A short history (I)

We need to compress losslessly – what options do we have?

- **Bit stream compressors: ZSTD on fast – 60% of real time (RT), 60% compression**
- **JPEG lossless: 8 bpp only**
- **H264: >10 bpp only with expensive commercial solutions (not tested)**

Reference: Intel Xeon, 3.4GHz, 24 cores; 128GB RAM, Geforce GTX 2060 6GB

X% compression ⇔ compressed size is X% of raw size

data from 2020

Fusion Power Video – A short history (I)

We need to compress losslessly – what options do we have?

- **Bit stream compressors: ZSTD on fast – 60% of real time (RT), 60% compression**
- **JPEG lossless: 8 bpp only**
- **H264: >10 bpp only with expensive commercial solutions (not tested)**
- **H265/VP9: (max. 12 bpp) – 550% of RT, 25% compression**

Reference: Intel Xeon, 3.4GHz, 24 cores; 128GB RAM, Geforce GTX 2060 6GB
X% compression ⇔ compressed size is X% of raw size
data from 2020

Fusion Power Video – A short history (I)

We need to compress losslessly – what options do we have?

- **Bit stream compressors: ZSTD on fast – 60% of real time (RT), 60% compression**
- **JPEG lossless: 8 bpp only**
- **H264: >10 bpp only with expensive commercial solutions (not tested)**
- **H265/VP9: (max. 12 bpp) – 550% of RT, 25% compression**
- **FFV1: 300% of RT, 50% compression**
- **FFVHUFF/Dirac: < 90% of RT, >90% compression …**

Reference: Intel Xeon, 3.4GHz, 24 cores; 128GB RAM, Geforce GTX 2060 6GB

X% compression ⇔ compressed size is X% of raw size

data from 2020

Fusion Power Video – A short history (I)

We need to compress losslessly – what options do we have?

- **Bit stream compressors: ZSTD on fast – 60% of real time (RT), 60% compression**
- **JPEG lossless: 8 bpp only**
- **H264: >10 bpp only with expensive commercial solutions (not tested)**
- **H265/VP9: (max. 12 bpp) – 550% of RT, 25% compression**
- **FFV1: 300% of RT, 50% compression**
- **FFVHUFF/Dirac: < 90% of RT, >90% compression …**

What about HW support? (Nvidia CUDA – Intel QSV lossy only)

Reference: Intel Xeon, 3.4GHz, 24 cores; 128GB RAM, Geforce GTX 2060 6GB

X% compression ⇔ compressed size is X% of raw size

data from 2020

Fusion Power Video – A short history (I)

We need to compress losslessly – what options do we have?

- **Bit stream compressors: ZSTD on fast – 60% of real time (RT), 60% compression**
- **JPEG lossless: 8 bpp only**
- **H264: >10 bpp only with expensive commercial solutions (not tested)**
- **H265/VP9: (max. 12 bpp) – 550% of RT, 25% compression**
- **FFV1: 300% of RT, 50% compression**
- **FFVHUFF/Dirac: < 90% of RT, >90% compression …**

What about HW support? (Nvidia CUDA – Intel QSV lossy only)

- **JPEG2000 CUDA: ~100% of RT, 77% compression (commercial codec, expensive)**

Reference: Intel Xeon, 3.4GHz, 24 cores; 128GB RAM, Geforce GTX 2060 6GB
X% compression ⇔ compressed size is X% of raw size
data from 2020

Fusion Power Video – A short history (I)

We need to compress losslessly – what options do we have?

- **Bit stream compressors: ZSTD on fast – 60% of real time (RT), 60% compression**
- **JPEG lossless: 8 bpp only**
- **H264: >10 bpp only with expensive commercial solutions (not tested)**
- **H265/VP9: (max. 12 bpp) – 550% of RT, 25% compression**
- **FFV1: 300% of RT, 50% compression**
- **FFVHUFF/Dirac: < 90% of RT, >90% compression …**

What about HW support? (Nvidia CUDA – Intel QSV lossy only)

- **JPEG2000 CUDA: ~100% of RT, 77% compression (commercial codec, expensive)**
- **H264/265 CUDA: only <= 10bpp**

Reference: Intel Xeon, 3.4GHz, 24 cores; 128GB RAM, Geforce GTX 2060 6GB

X% compression ⇔ compressed size is X% of raw size

data from 2020

Fusion Power Video – A short history (I)

We need to compress losslessly – what options do we have?

- **Bit stream compressors: ZSTD on fast – 60% of real time (RT), 60% compression**
- **JPEG lossless: 8 bpp only**
- **H264: >10 bpp only with expensive commercial solutions (not tested)**
- **H265/VP9: (max. 12 bpp) – 550% of RT, 25% compression**
- **FFV1: 300% of RT, 50% compression**
- **FFVHUFF/Dirac: < 90% of RT, >90% compression …**

What about HW support? (Nvidia CUDA – Intel QSV lossy only)

- **JPEG2000 CUDA: ~100% of RT, 77% compression (commercial codec, expensive)**
- **H264/265 CUDA: only <= 10bpp**
  - **8/10 MSB in Y, rest in U plane: 80% of RT, 36% compression**

Reference: Intel Xeon, 3.4GHz, 24 cores; 128GB RAM, Geforce GTX 2060 6GB
X% compression ⇔ compressed size is X% of raw size
data from 2020

Fusion Power Video – A short history (I)

We need to compress losslessly – what options do we have?

- **Bit stream compressors: ZSTD on fast – 60% of real time (RT), 60% compression**
- **JPEG lossless: 8 bpp only**
- **H264: >10 bpp only with expensive commercial solutions (not tested)**
- **H265/VP9: (max. 12 bpp) – 550% of RT, 25% compression**
- **FFV1: 300% of RT, 50% compression**
- **FFVHUFF/Dirac: < 90% of RT, >90% compression …**

What about HW support? (Nvidia CUDA – Intel QSV lossy only)

- **JPEG2000 CUDA: ~100% of RT, 77% compression (commercial codec, expensive)**
- **H264/265 CUDA: only <= 10bpp**
  - **8/10 MSB in Y, rest in U plane: 80% of RT, 36% compression**
  - **8/10 MSB in grayscale, rest in separate stream (bit packing): 60% of RT, 33% compression (~6% MSB)**

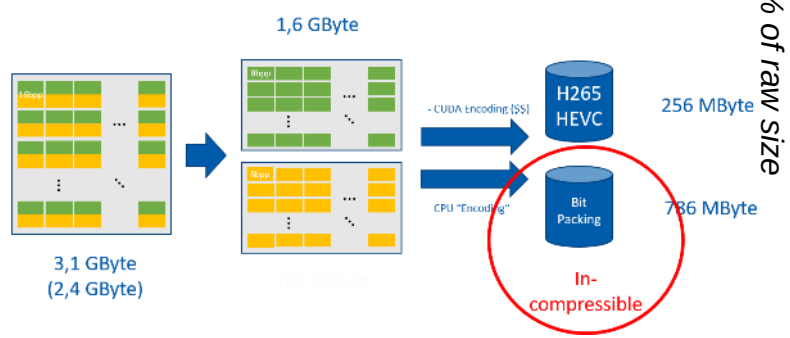

Reference: Intel Xeon, 3.4GHz, 24 cores; 128GB RAM, Geforce GTX 2060 6GB
X% compression ⇔ compressed size is X% of raw size
data from 2020

Fusion Power Video – A short history (II)

New upcoming, promising standard: JPEG XL

Fusion Power Video – A short history (II)

New upcoming, promising standard: JPEG XL

pro:

- **Support for lossless & 16+bpp**
- **Moderate compute resources**
- **(upcoming) open standard & open source**

Fusion Power Video – A short history (II)

New upcoming, promising standard: JPEG XL

pro:

- **Support for lossless & 16+bpp**
- **Moderate compute resources**
- **(upcoming) open standard & open source**

con:

- **not yet finalized (2020)**
- **Reference implementation not ready for reference testing**

Fusion Power Video – A short history (II)

New upcoming, promising standard: JPEG XL

pro:
- Support for lossless & 16+bpp
- Moderate compute resources
- (upcoming) open standard & open source

con:
- not yet finalized (2020)
- Reference implementation not ready for reference testing

Contacts to JPEG XL developers lead to cooperation with Google Research, Zurich

Google Research

Fusion Power Video – A short history (II)

New upcoming, promising standard: JPEG XL

pro:
- Support for lossless & 16+bpp
- Moderate compute resources
- (upcoming) open standard & open source

con:
- not yet finalized (2020)
- Reference implementation not ready for reference testing

Contacts to JPEG XL developers lead to cooperation with Google Research, Zurich

- Supported with in depth domain knowledge

Google Research

Fusion Power Video – A short history (II)

New upcoming, promising standard: JPEG XL

pro:
- Support for lossless & 16+bpp
- Moderate compute resources
- (upcoming) open standard & open source

con:
- not yet finalized (2020)
- Reference implementation not ready for reference testing

Contacts to JPEG XL developers lead to cooperation with Google Research, Zurich

- Supported with in depth domain knowledge
- Analyzed samples of real world W7-X videos

Google Research

Fusion Power Video – A short history (II)

New upcoming, promising standard: JPEG XL

pro:
- Support for lossless & 16+bpp
- Moderate compute resources
- (upcoming) open standard & open source

con:
- not yet finalized (2020)
- Reference implementation not ready for reference testing

Contacts to JPEG XL developers lead to cooperation with Google Research, Zurich

- Supported with in depth domain knowledge
- Analyzed samples of real world W7-X videos
- Identified, implemented & open sourced low compute, high impact W7-X specific compression scheme

Google Research

Fusion Power Video – A short history (II)

New upcoming, promising standard: JPEG XL

pro:
- Support for lossless & 16+bpp
- Moderate compute resources
- (upcoming) open standard & open source

con:
- not yet finalized (2020)
- Reference implementation not ready for reference testing

Contacts to JPEG XL developers lead to cooperation with Google Research, Zurich

- Supported with in depth domain knowledge
- Analyzed samples of real world W7-X videos
- Identified, implemented & open sourced low compute, high impact W7-X specific compression scheme
- Collaboratively evolved the C++ implementation

Google Research

Fusion Power Video – A short history (II)

New upcoming, promising standard: JPEG XL

pro:
- Support for lossless & 16+bpp
- Moderate compute resources
- (upcoming) open standard & open source

con:
- not yet finalized (2020)
- Reference implementation not ready for reference testing

Contacts to JPEG XL developers lead to cooperation with Google Research, Zurich

- Supported with in depth domain knowledge
- Analyzed samples of real world W7-X videos
- Identified, implemented & open sourced low compute, high impact W7-X specific compression scheme
- Collaboratively evolved the C++ implementation

Google Research

https://github.com/google/fusion-power-video

Fusion Power Video – Algorithm Details

Digital images | Compression | Challenges | S.o.t. Art | FPV History | **FPV Algorithm** | FPV@W7X | Outlook | Standardization

Fusion Power Video – Algorithm Details

Raw input image

Fusion Power Video – Algorithm Details

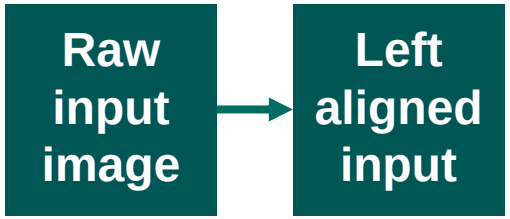

Left Shift to Alignment

Digital images | Compression | Challenges | S.o.t. Art | FPV History | **FPV Algorithm** | FPV@W7X | Outlook | Standardization

Fusion Power Video – Algorithm Details

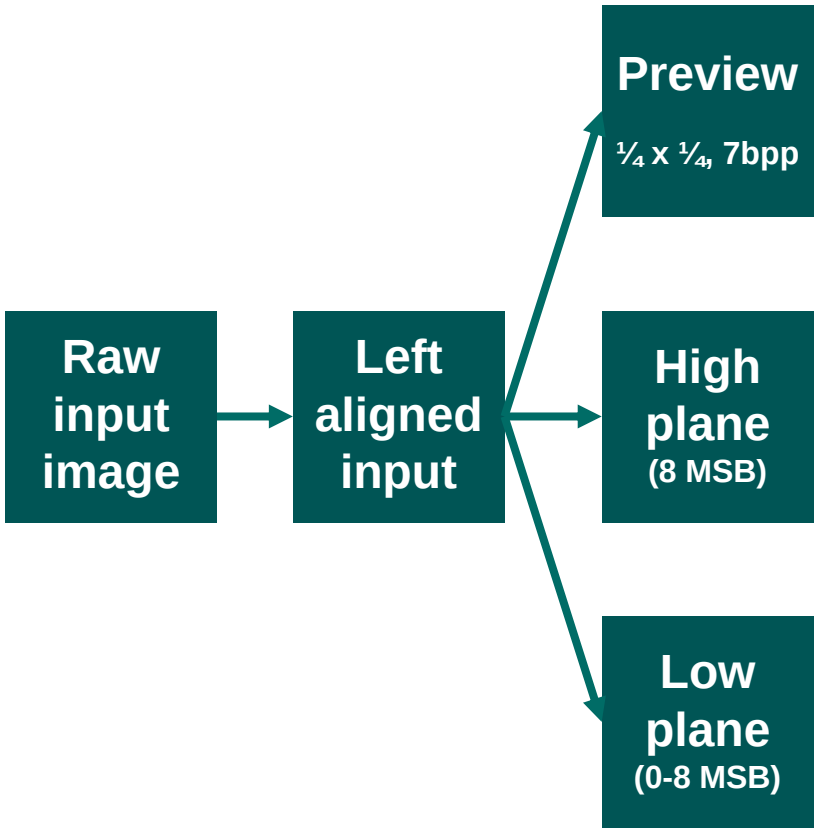

Digital images | Compression | Challenges | S.o.t. Art | FPV History | **FPV Algorithm** | FPV@W7X | Outlook | Standardization

Fusion Power Video – Algorithm Details

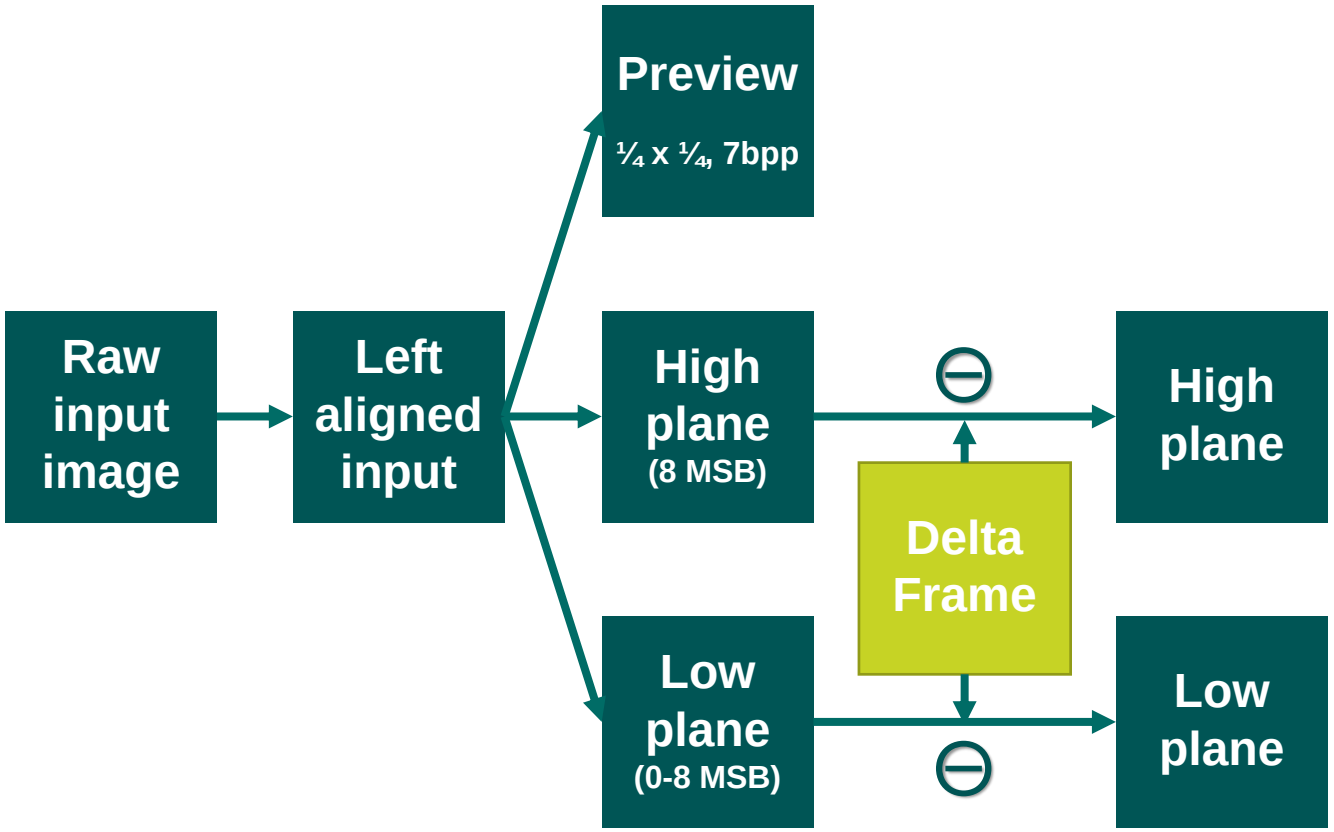

Digital images | Compression | Challenges | S.o.t. Art | FPV History | **FPV Algorithm** | FPV@W7X | Outlook | Standardization

Fusion Power Video – Algorithm Details

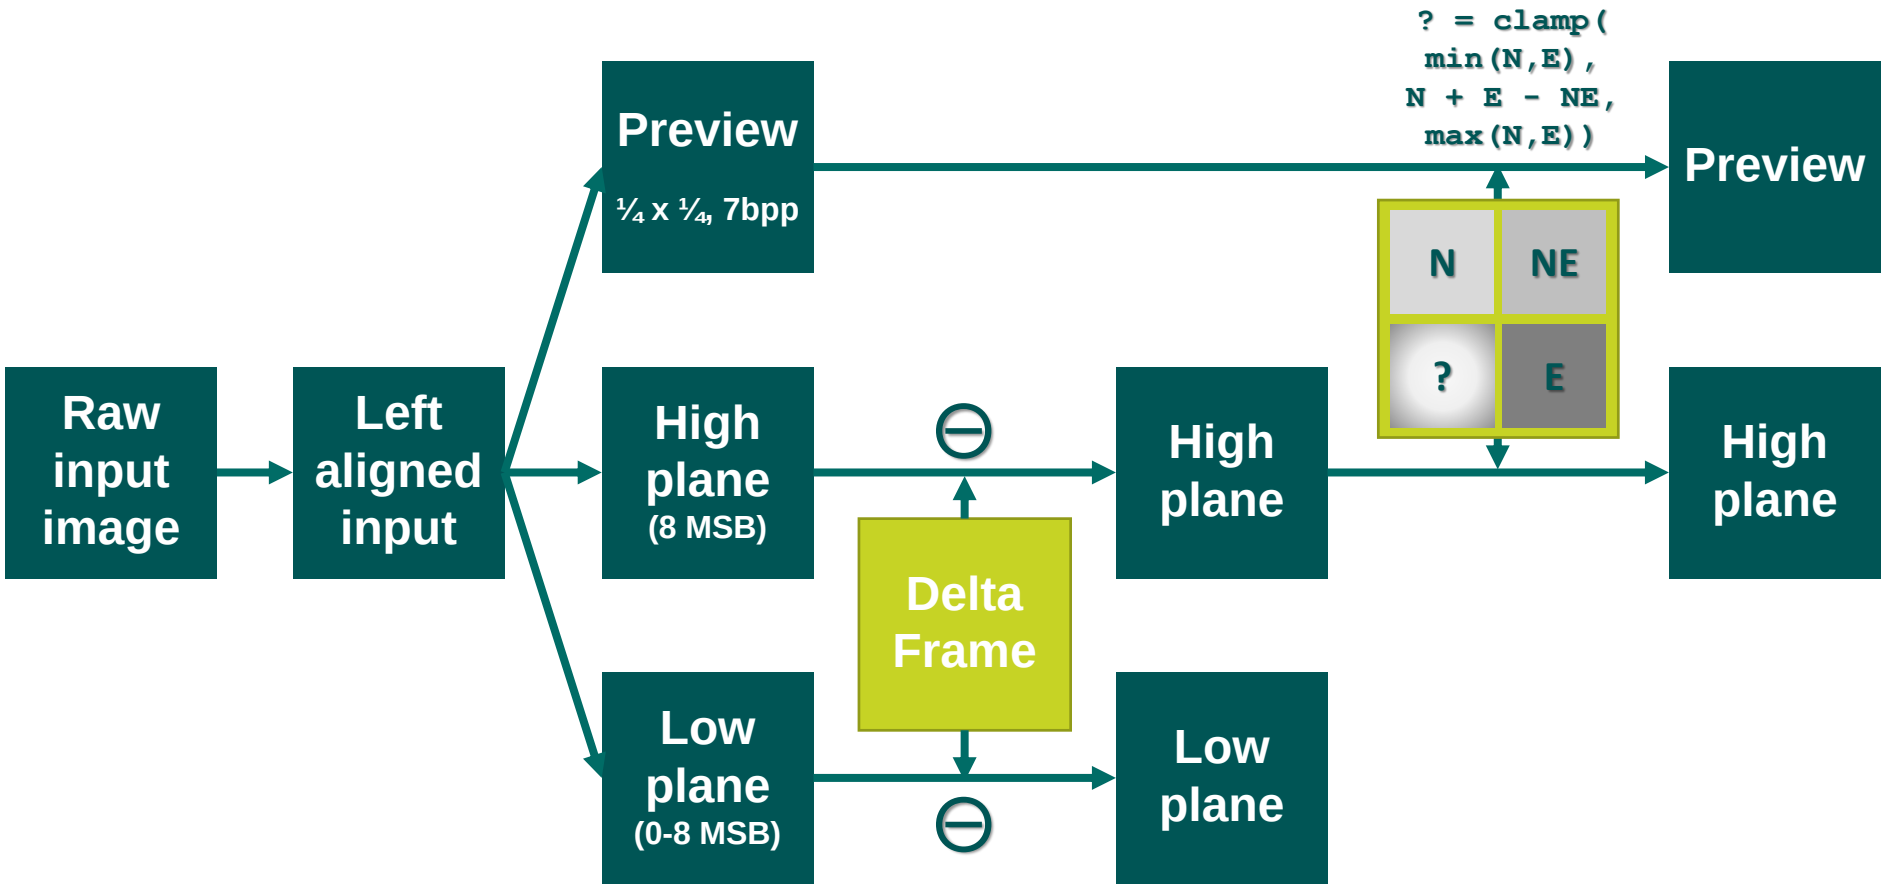

Left Shift to Alignment | **Plane Splitting** | **Delta Prediction** | **Clamped Gradient Prediction**

Digital images | Compression | Challenges | S.o.t. Art | FPV History | **FPV Algorithm** | FPV@W7X | Outlook | Standardization

Fusion Power Video – Algorithm Details

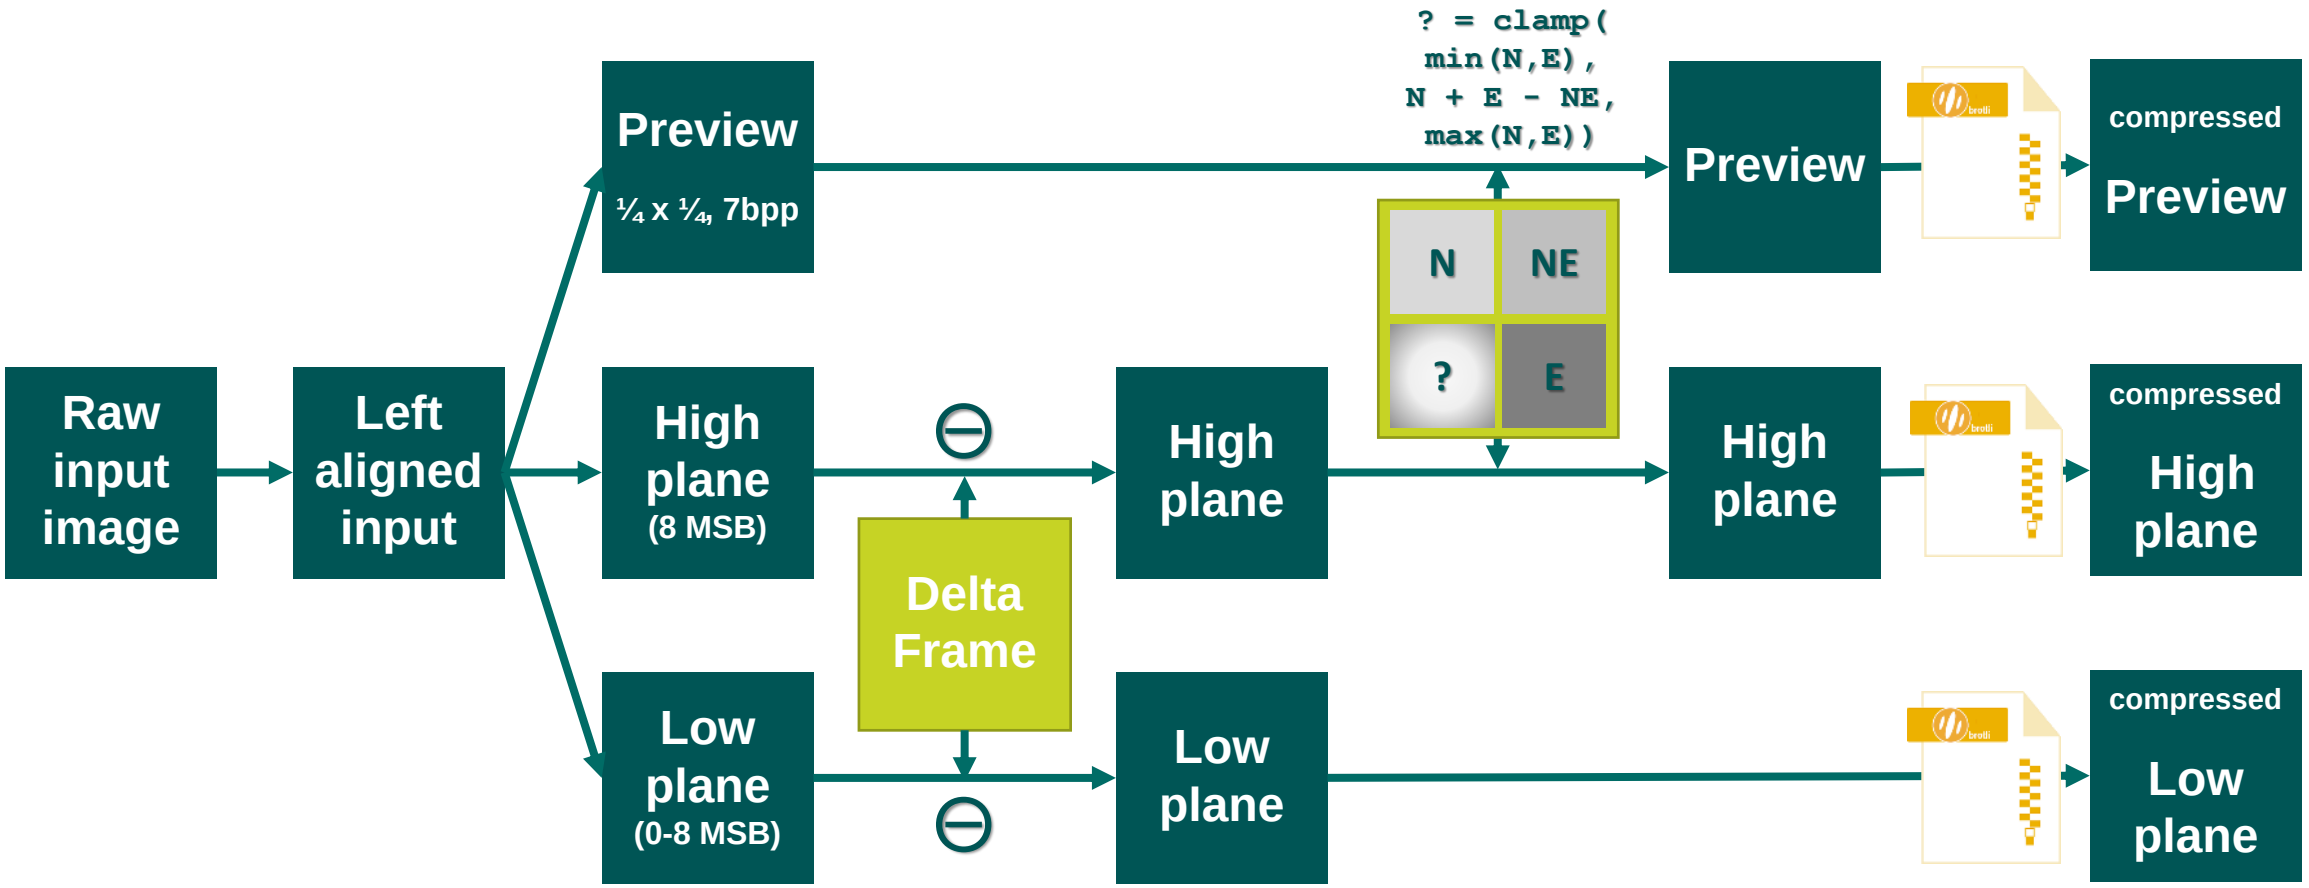

Left Shift to Alignment | **Plane Splitting** | **Delta Prediction** | **Clamped Gradient Prediction** | **brotli Stream Compression**

Fusion Power Video – Implementation at Wendelstein 7-X

Where did we go from there?

Digital images | Compression | Challenges | S.o.t. Art | FPV History | FPV Algorithm | **FPV@W7X** | Outlook | Standardization

Fusion Power Video – Implementation at Wendelstein 7-X

Where did we go from there?

- **Java wrapper around the C++ library**
  - **Performance issues at JNI boundary: copy ops & incompatible asynchrony frameworks**

Fusion Power Video – Implementation at Wendelstein 7-X

Where did we go from there?

- **Java wrapper around the C++ library**
  - **Performance issues at JNI boundary: copy ops & incompatible asynchrony frameworks**
- **Full reimplementation in native Java**
  - **Exception: brotli compression (uses brotli4j JNA wrapper of libbrotli)**

Fusion Power Video – Implementation at Wendelstein 7-X

Where did we go from there?

- **Java wrapper around the C++ library**
  - **Performance issues at JNI boundary: copy ops & incompatible asynchrony frameworks**
- **Full reimplementation in native Java**
  - **Exception: brotli compression (uses brotli4j JNA wrapper of libbrotli)**
- **Columnar serialization format**
  - **Batches of e.g. 100 frames**
  - **non-interlaced planes for efficient extraction of details**

Fusion Power Video – Implementation at Wendelstein 7-X

How to work with FPV data from W7-X ArchiveDB?

Digital images | Compression | Challenges | S.o.t. Art | FPV History | FPV Algorithm | **FPV@W7X** | Outlook | Standardization

Fusion Power Video – Implementation at Wendelstein 7-X

How to work with FPV data from W7-X ArchiveDB?

- **Decompression is fully transparent to the reader**

Fusion Power Video – Implementation at Wendelstein 7-X

How to work with FPV data from W7-X ArchiveDB?

- **Decompression is fully transparent to the reader**
- **SignalAccess and WebAPI offer 4 separate signals:**

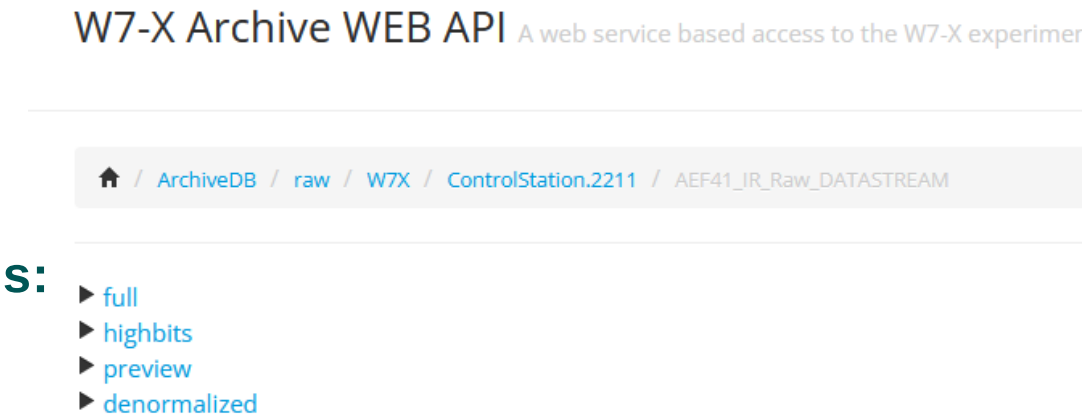

Digital images | Compression | Challenges | S.o.t. Art | FPV History | FPV Algorithm | **FPV@W7X** | Outlook | Standardization

Fusion Power Video – Implementation at Wendelstein 7-X

How to work with FPV data from W7-X ArchiveDB?

- Decompression is fully transparent to the reader
- SignalAccess and WebAPI offer 4 separate signals:
  - full: left aligned 16 bpp (values 0 – 65535)

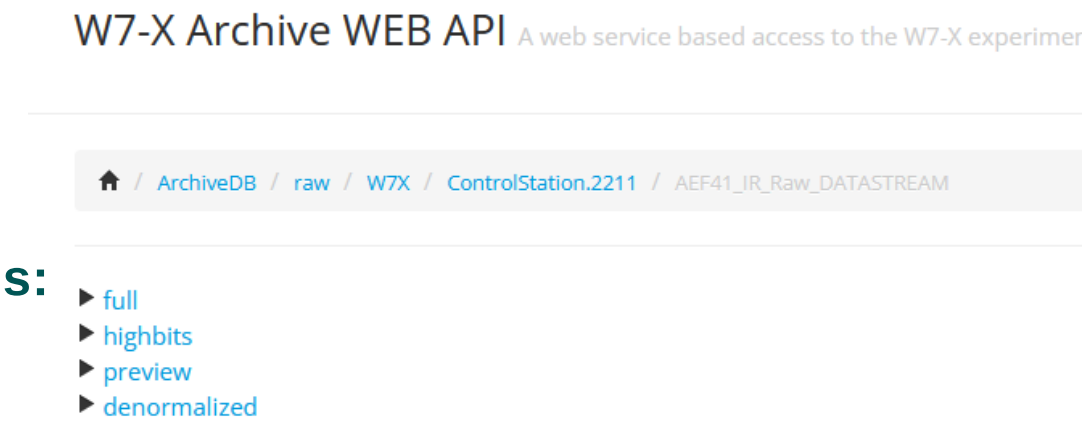

Digital images | Compression | Challenges | S.o.t. Art | FPV History | FPV Algorithm | **FPV@W7X** | Outlook | Standardization

Fusion Power Video – Implementation at Wendelstein 7-X

How to work with FPV data from W7-X ArchiveDB?

- Decompression is fully transparent to the reader
- SignalAccess and WebAPI offer 4 separate signals:
  - full: left aligned 16 bpp (values 0 – 65535)
  - highbits: left aligned 8 bpp MSBs (values 0 – 255)

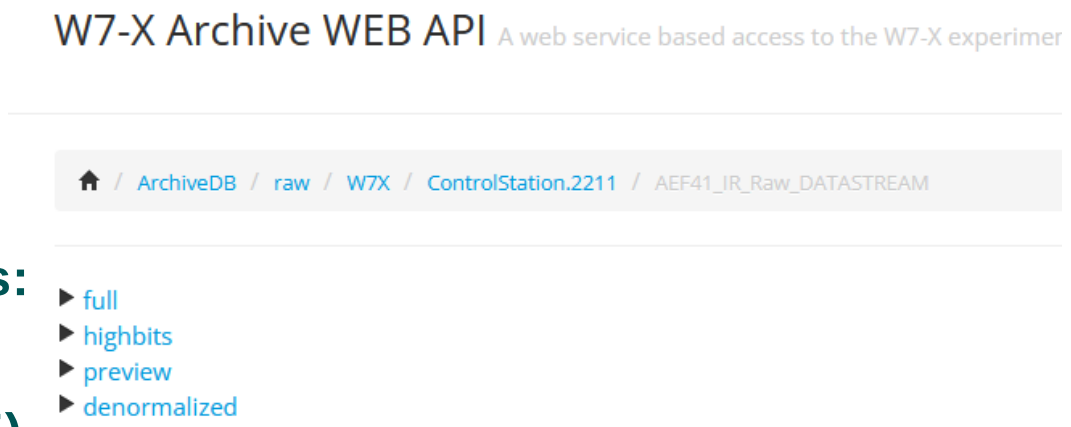

Digital images | Compression | Challenges | S.o.t. Art | FPV History | FPV Algorithm | **FPV@W7X** | Outlook | Standardization

Fusion Power Video – Implementation at Wendelstein 7-X

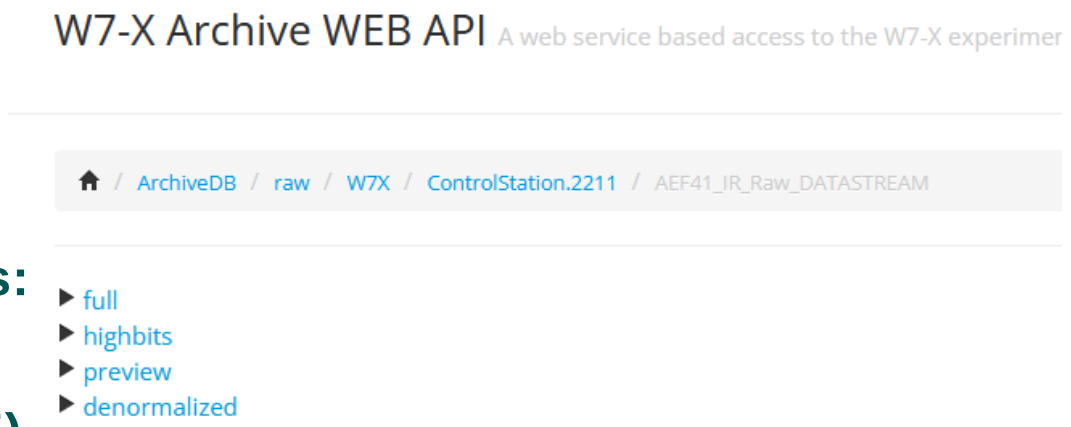

How to work with FPV data from W7-X ArchiveDB?

- **Decompression is fully transparent to the reader**
- **SignalAccess and WebAPI offer 4 separate signals:**
  - **full: left aligned 16 bpp (values 0 – 65535)**
  - **highbits: left aligned 8 bpp MSBs (values 0 – 255)**
  - **preview: left aligned 7 bpp, 4x4 subsampled (values 0 – 254)**

Digital images | Compression | Challenges | S.o.t. Art | FPV History | FPV Algorithm | **FPV@W7X** | Outlook | Standardization

Fusion Power Video – Implementation at Wendelstein 7-X

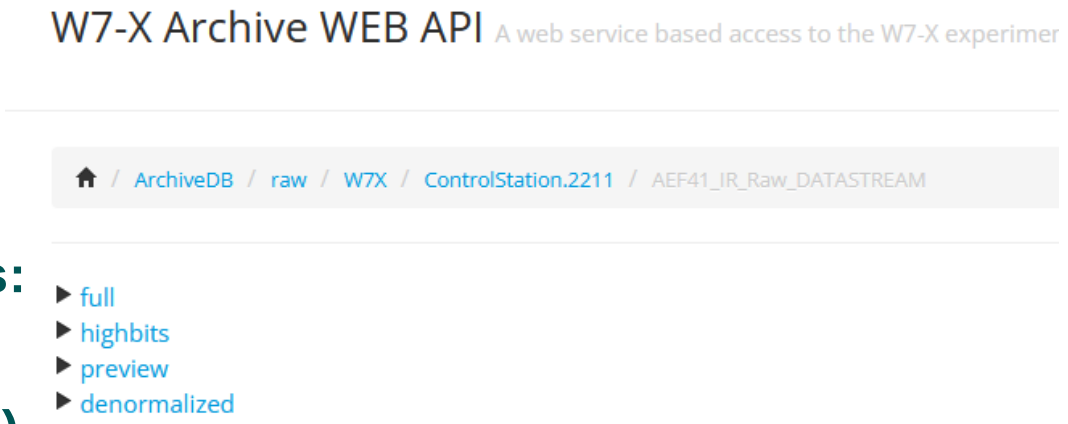

How to work with FPV data from W7-X ArchiveDB?

- **Decompression is fully transparent to the reader**
- **SignalAccess and WebAPI offer 4 separate signals:**
  - **full: left aligned 16 bpp (values 0 – 65535)**
  - **highbits: left aligned 8 bpp MSBs (values 0 – 255)**
  - **preview: left aligned 7 bpp, 4x4 subsampled (values 0 – 254)**
  - **denormalized: reconstructed original alignment, 16 bpp (values camera dependent)**

Fusion Power Video – Implementation at Wendelstein 7-X

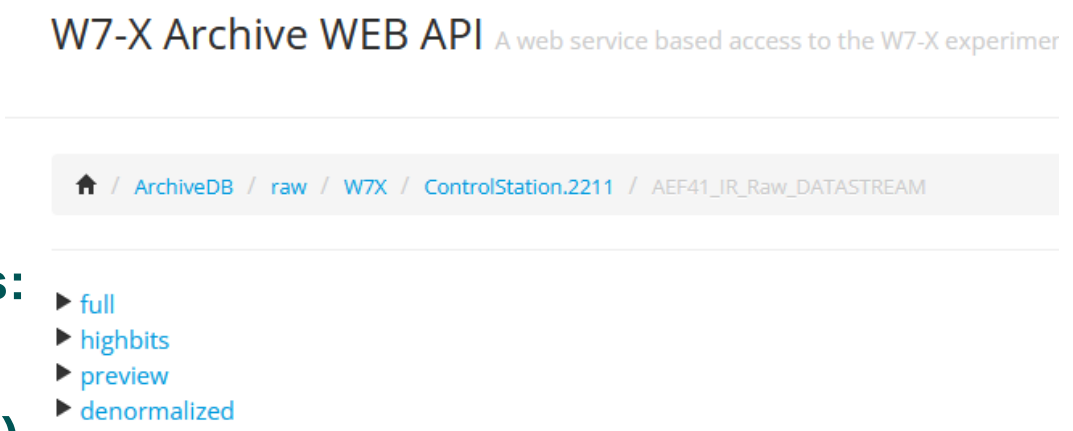

How to work with FPV data from W7-X ArchiveDB?

- Decompression is fully transparent to the reader
- SignalAccess and WebAPI offer 4 separate signals:
  - full: left aligned 16 bpp (values 0 – 65535)
  - highbits: left aligned 8 bpp MSBs (values 0 – 255)
  - preview: left aligned 7 bpp, 4x4 subsampled (values 0 – 254)
  - denormalized: reconstructed original alignment, 16 bpp (values camera dependent)

- Current limitations:
  - Client side decompression using Java API (SignalAccess) – including full data download even for preview

Fusion Power Video – Implementation at Wendelstein 7-X

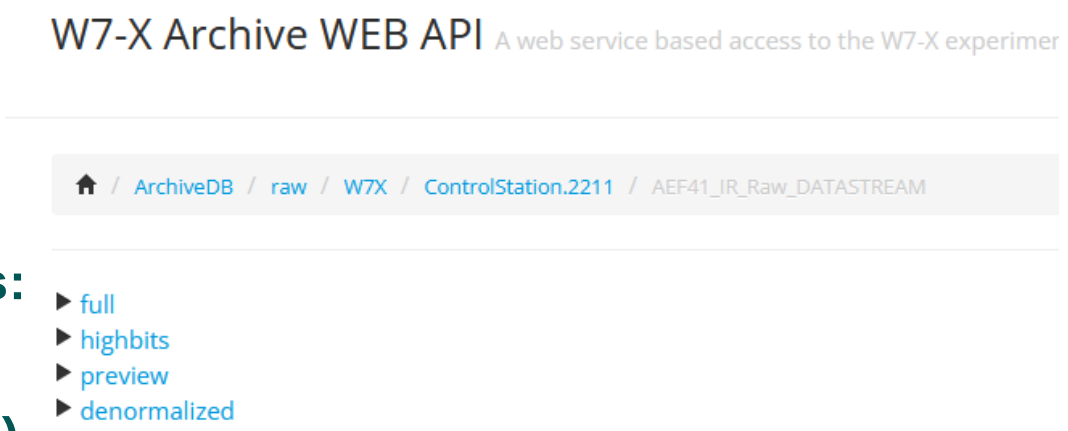

How to work with FPV data from W7-X ArchiveDB?

- Decompression is fully transparent to the reader
- SignalAccess and WebAPI offer 4 separate signals:
  - full: left aligned 16 bpp (values 0 – 65535)
  - highbits: left aligned 8 bpp MSBs (values 0 – 255)
  - preview: left aligned 7 bpp, 4x4 subsampled (values 0 – 254)
  - denormalized: reconstructed original alignment, 16 bpp (values camera dependent)

- Current limitations:
  - Client side decompression using Java API (SignalAccess) – including full data download even for preview
  - ***With compression, we pay space savings with compute: reading is slower than on raw data***

Outlook

Outlook

- ***Try to go open source***

Outlook

- ***Try to go open source***
- **Implement read optimizations for preview & high planes**

Outlook

- ***Try to go open source***
- **Implement read optimizations for preview & high planes**
- **Prevent delta frame duplication**

Outlook

- ***Try to go open source***
- **Implement read optimizations for preview & high planes**
- **Prevent delta frame duplication**
- **Directly use FPV preview in W7-X DataMonitor**

Outlook

- ***Try to go open source***
- **Implement read optimizations for preview & high planes**
- **Prevent delta frame duplication**
- **Directly use FPV preview in W7-X DataMonitor**
- **Allow WebAPI Archive uploads with compression / of compressed data**

Outlook

- ***Try to go open source***
- **Implement read optimizations for preview & high planes**
- **Prevent delta frame duplication**
- **Directly use FPV preview in W7-X DataMonitor**
- **Allow WebAPI Archive uploads with compression / of compressed data**

- **Since it is standardized these days: evaluate jpegXL**
  - **Would offer colored images, maybe multispectral, maybe float images, more generalized image compression**
  - **Compare performance of FPV preview and progressive jpegXL**
  - **Compare performance of FPV high byte only extraction and progressive jpegXL**
  - **maybe combine approaches for future jpegXL based lossless video compression scheme**

From first principles – a call for standardization

Digital images | Compression | Challenges | S.o.t. Art | FPV History | FPV Algorithm | FPV@W7X | Outlook | **Standardization**

From first principles – a call for standardization

- ***Let all be aligned to the left!***
- ***Don’t stuff your tummy with waste!***
- ***All power to the little people!***
- ***All types are not created equal!***

Digital images | Compression | Challenges | S.o.t. Art | FPV History | FPV Algorithm | FPV@W7X | Outlook | **Standardization**

From first principles – a call for standardization

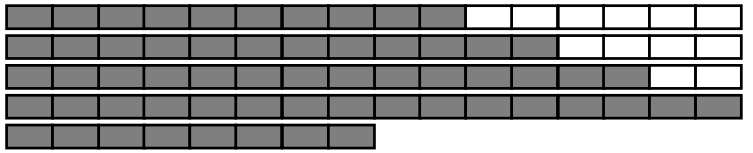

- ***Let all be aligned to the left!***
  - **When working with relative intensity images, normalize everything to 65,535**
- ***Don’t stuff your tummy with waste!***
- ***All power to the little people!***
- ***All types are not created equal!***

From first principles – a call for standardization

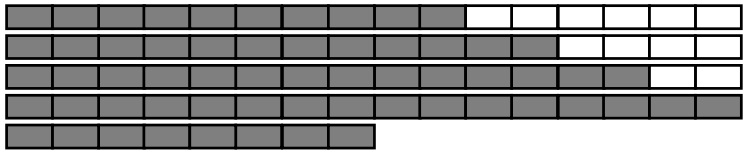

- ***Let all be aligned to the left!***
  - **When working with relative intensity images, normalize everything to 65,535**

- ***Don't stuff your tummy with waste!***
  - **Before compressing images, remove what's of no value**

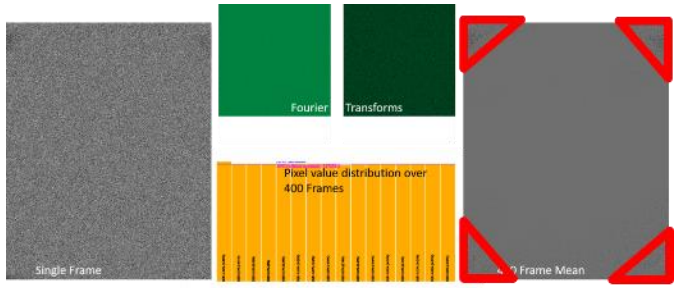

- ***All power to the little people!***

- ***All types are not created equal!***

From first principles – a call for standardization

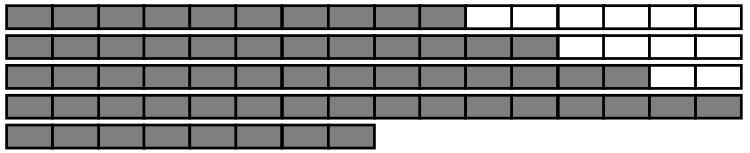

- ***Let all be aligned to the left!***
  - When working with relative intensity images, normalize everything to 65,535

- ***Don't stuff your tummy with waste!***
  - Before compressing images, remove what's of no value

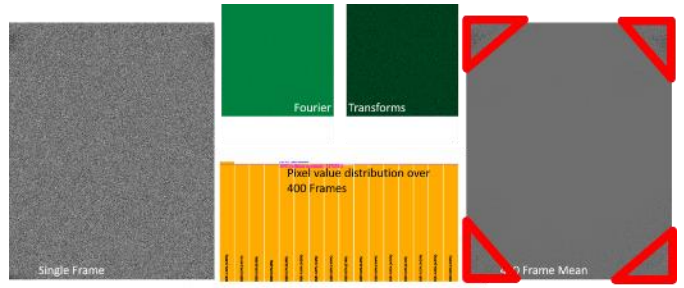

- ***All power to the little people!***
  - Almost all compute hardware today (for integers) is LITTLE ENDIAN – use it

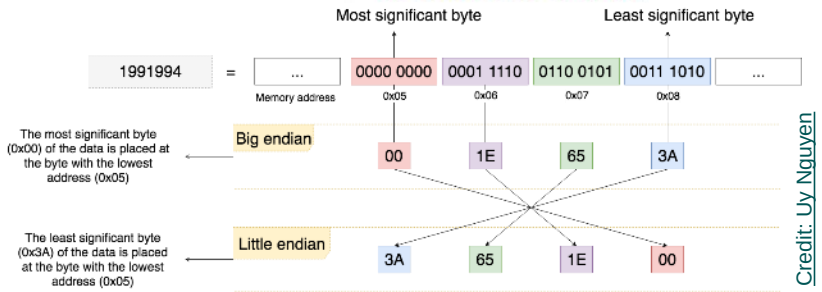

- ***All types are not created equal!***

From first principles – a call for standardization

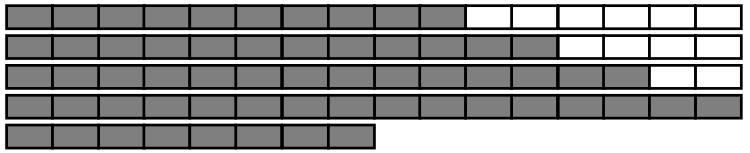

- ***Let all be aligned to the left!***
  - When working with relative intensity images, normalize everything to 65,535

- ***Don't stuff your tummy with waste!***
  - Before compressing images, remove what's of no value

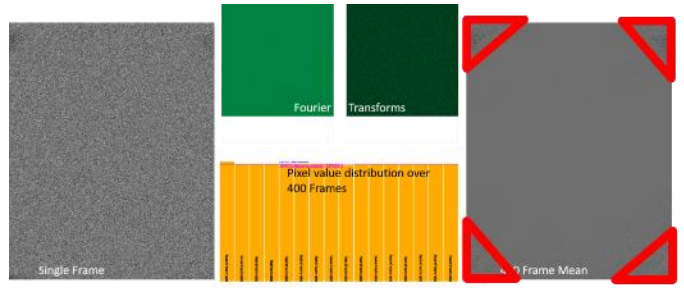

- ***All power to the little people!***
  - Almost all compute hardware today (for integers) is LITTLE ENDIAN – use it

- ***All types are not created equal!***
  - Generally, IEEE floating point does more harm than help – prefer fixed point == integer

```
0.1 + 0.2 != 0.3
0/0 != 0/0
-0 == +0
```

Questions?

Thank you

very much

for your attention!

Simon Fischer
simon.fischer@ipp.mpg.de
@simon-fischer-24255a109
@percurious

Basics – Compression schemes – predictive encoding

Predict the next value based on previous one(s)

4	4	5	6	5	5	3	4	4	3	4	5
13	14	15	16	16	17	18	20	20	21	22	21
4	42	23	6	15	-7	8	9	9	27	-4	5

- Prediction: Each value is equal to its predecessor

Subtract the predicted from the actual value

4	0	1	1	-1	0	-2	1	0	-1	1	1
13	1	1	1	0	1	1	2	0	1	1	-1
4	38	-19	-17	9	-22	15	1	0	18	-31	9

- Good prediction: values of/around 0 are dominant
- Bad prediction: values vary a lot

⬇

Good prediction: Variable length encoding (maybe even RLE) is very effective

(bad prediction ➲ encoding effect low/negative)

In images, predictions can utilize spatial coherence – in video also temporal