SKA1 SDP Risk Register

Document Number..... SKA-TEL-SDP-0000052
Document Type..... REF
Revision..... 08
Authors..... F. Graser, P. Alexander, J. Coles, V. Allan
Date..... 2018-10-31
Document Classification.....FOR PROJECT USE ONLY
Status..... Released

Name	Designation	Affiliation	Signature	
Authored by:				
Ferdl Graser	SDP Systems Engineer	SARAO	Ferdl Graser (Oct 31, 2018)	
			Date:	
Owned by:				
			Date:	
Approved by:				
			Date:	
Released by:				
Paul Alexander	SDP Project Lead	University of Cambridge	Paul Alexander (Oct 31, 2018)	
			Date:	

DOCUMENT HISTORY

Revision	Date Of Issue	Engineering Change Number	Comments
01	2015-02-2015		Prepared for M7, inputs to System PDR
03	2016-01-26		Prepared for SDP PDR
04			SDP PDR Baseline
04B			Prepared for M12 SDP design for Prototyping
04C	2016-04-08		Prepared for M14, SDP delta-PDR
05	2016-08-08		Prepared for M15 Risk Review 1
06			Prepared for M17 Risk Review 3
07			Prepared for M18 SDP Risk Review 4
07C		ECP-160048 ECP-160057 ECP-180001	SDP L1s related to the SKA!-LOW EoR pipeline SKA Observatory Systems: PBS re-alignment prepared for M20
08	2018-10-31	ECP-SDP-180002	Updated for SDP CDR submission (M21) SDP PRECDRs: SDPPRECDR-5 “Document seems not to be used” SDPPRECDR-54 ‘‘SDPRISK-359’’ SDPPRECDR-55 ‘‘Risk Unclear’’ SDPPRECDR-58 ‘‘Risk responsibility unclear’’ SDPPRECDR-59 ‘‘Are contingency included for these risks?’’ SDPPRECDR-62 ‘‘SDPRISK-374’’ SDPPRECDR-63 ‘‘SDPRISK-413’’ SDPPRECDR-65 ‘‘SDPRISK-313’’ SDPPRECDR-66 ‘‘SDPRISK-363 - Buffer hardware and software does not meet performance requirements’’ SDPPRECDR-67 ‘‘SDPRISK-392’’ SDPPRECDR-68 ‘‘SDPRISK-397’’ SDPPPRECDR-69 ‘‘SDPRISK-408’’ SDPPRECDR-163 ‘‘SDPRISK-343’’ SDPPRECDR-307 ‘‘SDPRISK-397 - Power model underestimates power consumption’’ SDPPRECDR-308 ‘‘SDPRISK-315 - Inadequate System Performance due to the interface between the Data Processor and Local Telescope State’’ SDPPRECDR-310 ‘‘SDPRISK-358: SDP extension funding reduced - Impact: late or reduced quality deliverables’’ SDPPRECDR-313 ‘‘SDPRISK-401 - Reduction of science performances’’ SDPPRECDR-64 ‘‘compute islands’’ SDPPRECDR-164 ‘‘SDPRISK-320, -341, -315 and others’’ SDPPRECDR-171 ‘‘Failure modes’’ SDPPRECDR-305 ‘‘The risk register does not include an explanation of the meaning of the risk levels’’

DOCUMENT SOFTWARE

	Package	Version	Filename
Word processor	Google Docs		SKA-TEL-SKO-0000000-01_GenDocTemplate
Block diagrams			
Google docs Add-ons	Table of contents		Used for heading numbering.
Risk Register	JIRA, with Projectbalm Risk Register Addon		used to maintain the risk register, and to export the register.

ORGANISATION DETAILS

Name	SDP Consortium
Lead Organisation	The Chancellor, Masters and Scholars of the University of Cambridge The Old Schools Trinity Lane Cambridge CB1 ITN United Kingdom
Website	www.ska-sdp.org

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF ABBREVIATIONS

CA	Consortium Agreement
CDR	Critical Design Review
MoU	Memorandum of Understanding
PDR	Preliminary Design Review
RRB	Risk Review Board
SDP	Science Data Processor
SE	System Engineering
SKA	Square Kilometre Array (Organisation)
WBS	Work Breakdown Structure

1 Introduction

1.1 Purpose of the document

This document provides details of SDP risks that remain open going into bridging and construction. It provides the working basis and terminology upon which SDP risks have been categorised and managed.

1.2 Scope of the document

The document only covers identified open/live risks at the time of CDR submission. It does not attempt to provide a full history of SDP risks nor how they have directed SDP activities in SDP Sprint Planning.

2 References

2.1 Applicable documents

The following documents are applicable to the extent stated herein. In the event of conflict between the contents of the applicable documents and this document, **the applicable documents** shall take precedence.

AD01 SKA-MGT-RMT-MP-00001 SKA Risk Management Plan
AD05 SKA-TEL-SDP-0000013 SDP Architecture Rev 06

2.2 Reference documents

The following documents are referenced in this document. In the event of conflict between the contents of the referenced documents and this document, **this document** shall take precedence.

AD01 SKA-TEL-SDP-0000092 M15: SDP Development Risk Review 1
AD02 SKA-TEL-SDP-0000107 M16: SDP Development Risk Review 2
AD03 SKA-TEL-SDP-0000109 M17: SDP Development Risk Review 3
AD04 SKA-TEL-SDP-0000112 M18: SDP Development Risk Review 4

3 Summary

The SDP Consortium has actively managed risk using a risk register throughout the design phase. Initially this was spreadsheet based, but for the last 2 years a JIRA plugin has been used. Section 4 provides a JIRA export of SDP open risks modified to fit the current document format. The Consortium risk management approach and categorisation has followed the guidelines of the SKAO laid out in the SKA Risk Management Plan [AD01]. That document uses the following definition of risk:

"An uncertain event or set of events that, should it occur, will have an effect on the achievement of SKA Project objectives. A risk is measured by the combination of the probability of a perceived threat occurring and the magnitude of its impact on the attainment of objectives".

SDP risk scoring uses fields aligned with the SKA Risk Register [AD01]. To ease reading of this document a brief summary is given to explain how SDP risks have been scored and the terminology used in the risk register. At the time of CDR submission the SDP design is assessed to have no residual extreme risks, 2 high risks, 13 medium and 3 low. This document provides a top-level view of these remaining SDP risks. Throughout the SDP CDR documents reference is made to the unique IDs of risks using their JIRA identifiers of the form SDPRISK-ID where "ID" is the risk number. Risks have been raised, morphed, and closed throughout the project lifetime and therefore the size of the ID is unimportant but it is unique.

3.1 Key terms and risk scoring

SDP has relied on the following definitions of terms for the assessment of risks.

Risk Event
This is what we are working to avoid or reduce the likelihood or impact of occurring. Risks are future events that could interfere with achievement of objectives.

Risk Cause
This records the triggers, sources or circumstances that could act alone or together to increase the likelihood of the Risk Event occurring. There are usually multiple causes leading to a Risk Event.

After initial identification and ownership of risks, each risk has been assessed at regular intervals by a Risk Review Board (RRB). The initial assessment (without any mitigations applied) provide the risk impact/consequence and probability scores.

Impact/Consequence
If this Risk Event did occur, how would it impact objectives? What are the longer-term or cumulative consequences? The scoring system that has been applied is shown in Table 1. There were three main areas where SDP design and Consortium/Project risks were assessed: Cost; Schedule and Performance as defined in Table 2 (See [AD01] for the full explanations).

Consequence	Impact level	Cost	Schedule	Performance
Insignificant	1	Lower than 2% impact	Very minor or no slip in milestone (week)	Very minor or no impact
Minor	2	Impact between 2% & 5%	Minor slip in milestone (weeks)	Minor functional impact, or some reduction in performance; performance still acceptable
Moderate	3	Impact between 5% & 10%	Moderate slip in milestone (several weeks to 2 months)	Moderate functional impact or reduction in performance, performance marginally acceptable but may require minor design improvement
Major	4	Impact between 10% & 20%	Major slip to milestone (2 – 6 months)	Major functional impact or reduction in performance, performance not acceptable; requires design change
Catastrophic	5	Impact larger than 20%	Critical slip to milestone (more than 6 months)	Critical functional impact or reduction in performance, performance not acceptable; requires new design

Table 1: Consequence and impact levels

Cost	This covers direct and indirect consequences. The Cost impact is the cost increase due to the risk occurring compared to the current approved budget for the area of work affected.
Schedule	Time impact covers the possible delay caused by the risk to the following milestones: • 'Handover to Observatory Operations'
Performance	This relates to the impact of the risk on the ability of the equipment/infrastructure to perform as required and to specification.

Table 2: Definition of impact categories

Probability/Likelihood

The probability that a risk occurs has been scored according to the intervals given in Table 3.

Likelihood	Likelihood level	Probability of Occurrence
Not Likely	1	0 to 10%
Low Likelihood	2	11% to 39%
Likely	3	40% to 60%
Highly Likely	4	61% to 79%
Near Certainty	5	80% to 100%

Table 3:The probability intervals used to assess the risk likelihood.

The defined level of risk results from multiplying the probability score by the impact score. It is a value value between 1 and 25 (if any items are scored 0 they are removed as a risk). This leads to a categorisation of: Extreme; High; Medium and Low. As well as the initial risk scoring, regular residual assessments have been undertaken to provide a current view on the risk and these "residual" scorings give the current Consortium view on the risk remaining. In Figure 1 the residual SDP risks are shown with the categorisation colour coded (and number per category in brackets) as Red (0) Extreme; Orange (2) high; Yellow (3) medium and Green (3) low.

Figure 1: Open SDP risks at the time of CDR submission broken down into 4 risk categories.

3.2 Other risk fields and terms

The Risk Register extracts in section 4 of this document show additional fields that have helped with risk management. These are:

Treatment: The current "treatment" state of the risk. This is the strategy used to manage the risk and has the following states:

- Accept (do nothing and accept the possible consequences);
- Avoid (do not perform the activity causing the risk);
- Mitigate (devise a plan to reduce either the probability or the impact of the risk); and
- Transfer (make another party responsible for the risk).

Status: The current state of the risk. It could be Treated; Closed; Transferred or Avoided but only after the treatment mentioned above has been applied.

Treatment Plan: This details the approach to reducing the probability or impact of the risk.

3.3 SDP Development Risk Reviews

The present document sets out to give a factual representation of the current state of the SDP Risk Register. The SDP Consortium has been managing risk actively throughout the project and has presented a full analysis of the project and element risks at regular intervals in a series of milestone documents entitled "SDP Development Risk Review" [RD01, RD02, RD03, RD04]. Those documents give detailed accounts of how the Consortium has managed its program of work to reduce the level of risk exposure at CDR.

4 Risk register status

4.1 Context

Section 4.3 provides a direct export from the live SDP Risk Register which is held in the SDP JIRA system. Since it is a raw export more detail is provided than is strictly needed for CDR, however this does provide insight into the risk history, how opinions changed and linkages with other project management tools, as well as the SDP approach to managing risks. Whilst JIRA provides full explicit traceability, for this submission assignees and commentators have had their names removed. For CDR the Consortium is only including risks that remain open at the time of submission. At peak the Consortium was managing of order 130 risks spread across all categories.

In the following extracts additional fields are included for context. The **Project** is the JIRA area (since SDP also uses JIRA for task management we distinguished this area by creating a separate project), the **Component/s** are (where applicable) the originating area of the risk in the Product Breakdown Structure, **Type** is always "Risk" here (it could be Task etc.); **Reporter** is the individual or team responsible for managing the risk, the **Assignee** is the technical lead responsible for technical assessments, **Labels** gives a way of grouping risks by type (e.g. pre-construction) and **Issues Links** provides the association with SDP work tasking. Finally, the **Comments** sections have been included "as is" to show how risks have been actively managed and progressed.

Our recommendation for surveying these risks is to focus in turn on the Risk title, Risk Event, Risk Cause, Treatment Plan and Residual Exposure. Delving into the comments may be useful to give further context.

4.2 Open risk areas

Broadly the remaining identified SDP risks fall into the following categories:

Performance (320; 324; 336; 339; 344; 363;)
Implementation (392; 311; 390)
Design completeness (343)
Modelling uncertainties (360; 369; 401)
Requirements (346; 359; 360; 408).

We believe these risks are understood and at an acceptable level to proceed to (the bridging and) construction work of the SDP.

4.3 Remaining open risks

[SDPRISK-320] Fast imaging pipeline does not have the required performance to meet latency requirements. Created: 20/Sep/16 Updated: 19/Oct/18	
Status:	Treated
Project:	SDP Risk Register
Component/s:	C.1.2.2 Pipelines

Type:	Risk		
Reporter:	[27]	**Assignee:**	[1]
Labels:	Failure, Pre-Con, risks		

Issue Links:				
	Blocks			
	is blocked by	TSK-1904	Review risk 320	In Progress
Impact:	4 Major			
Probability:	4 Highly likely			
Exposure:	High			
Treatment:	Mitigate			
Treatment plan:	We recommend this work be taken forward in bridging to further mitigate - bring in the science teams to help define the workflow in a better way. And determine if the SDP limitations here are acceptable.			
Residual Impact:	4 Major			
Residual Probability:	3 Likely			
Residual Exposure:	High			
Risk Event:	System Performance is Impacted. Real time pipelines do not have sufficient performance to meet the latency requirement, data quality is reduced or science output is reduced.			
Risk Cause:	Uncertainty in the workflows that could be written for fast imaging means that SDP has insufficient performance against those particular workflows to meet the latency requirements.			

Comments

Comment by [2] [28/Sep/16]
We need further analysis to evaluate the risk. Therefore, we should not score it at this stage; we do not understand enough of the failure mode. We don't know how much data we'd lose, we don't know whether we'd lose that much information to the latency.Ferdl reviewed this with Ronald. This analysis task needs to be added to the WBS for the Fast Imaging Pipelines, and prioritised.
Comment by [3] (Inactive) [28/Sep/16]
I have added a task to the PIP WBS
Comment by [2] [04/Nov/16]
Bojan noted (23/06/16):5. The cause and event appear transposed in risk SDPR-M14-39.6. Risk SDPR-M14-39 should be clarified/elaborated. If I understand correctly, the risk event is that the SDP execution engine can not execute certain types of processing which are required for correct operation of the SDP. This will then require a redesign of the execution engine and potentially other components. The cause of the risks is that functional requirements on the execution engine have not been identified correctly.One example of issue of this category is a convergence-based stopping criterion for the

self-calibration loop. This requires the execution engine to make decisions on what to execute on basis of values of intermediate data products. It is currently unclear if this will be supported in the execution engine.In mitigation of this risk we list “functionality system integration prototype” which does not exist. I think we meant “functional execution prototype”?
Comment by [2] [08/Dec/16]
Failure mode: Fast telescope state meta data is not received in time to meet the fast preprocessing latency. This causes reduced data for fast imaging. Current design does not take failure mode into account.Bojan noted (23/06/16):5. The cause and event appear transposed in risk SDPR-M14-39.6. Risk SDPR-M14-39 should be clarified/elaborated. If I understand correctly, the risk event is that the SDP execution engine can not execute certain types of processing which are required for correct operation of the SDP. This will then require a redesign of the execution engine and potentially other components. The cause of the risks is that functional requirements on the execution engine have not been identified correctly.One example of issue of this category is a convergence-based stopping criterion for the self-calibration loop. This requires the execution engine to make decisions on what to execute on basis of values of intermediate data products. It is currently unclear if this will be supported in the execution engine.In mitigation of this risk we list “functionality system integration prototype” which does not exist. I think we meant “functional execution prototype”?
Comment by [2] [08/Dec/16]
Time score: 4.0
Comment by [2] [08/Dec/16]
Fin score: 2.0
Comment by [2] [08/Dec/16]
Safety score: 1.0
Comment by [2] [08/Dec/16]
Risk Number: 108.0
Comment by [2] [08/Dec/16]
Reputation score: 1.0
Comment by [2] [08/Dec/16]
performance score: 4.0
Comment by [2] [28/Feb/17]
Treatment: we have defined failure modes for the fast imaging design likelihood: likely We have retired a lot of the risk.
Comment by [2] [23/Oct/17]
This has a major impact on one requirement, but this isn't a major impact for the whole SDP system. (Note - some of our earlier analysis is not just about the fast imaging pipeline - it's related to data ingest instead.) We need to reanalyse this risk, with the [3] input - is this risk really about Fast Imaging, or is it another system risk? Also, if Fast Imaging goes wrong, can it degrade gracefully? (ie. can it run with reduced performance?) This should be part of a larger assessment on the risks of the real-time processing.
Comment by [4] (Inactive) [22/Nov/17]
[5] just to say that we even mentioned you (RRB21 23OCT17) as someone that can add some value to this matter working with Ronald
Comment by [2] [21/Mar/18]

Work was assigned, but was unable to be completed. See linked ticket. Cannot rescore.
Comment by [2] [26/Apr/18]
Want to take forward in 2018C.
Comment by [6] [04/Sep/18]
Assigned TSK-1904 to [1] to see if Mark can assess the impact of this risk.
Comment by [6] [04/Sep/18]
I've removed the **Treatment plan** text since this risk is not about failure modes, but rather about meeting the latency requirement.
Comment by [7] [19/Oct/18]
We could accept this risk as a limitation. We have L1 requirements in this area. We recommend this work be taken forward in bridging to further mitigate - bring in the science teams to help define the workflow in a better way. And determine if the SDP limitations here are acceptable.
Comment by [1] [19/Oct/18]
The prototyping work being led by [17] may help to inform us about the performance of the pipeline. The reports are here: https://v1.overleaf.com/read/pgwctzthqzds https://v1.overleaf.com/read/rzsrkbhfsvdp

[SDPRISK-392] Complexity of implementing very high performance processing components Created: 30/Jul/15 Updated: 19/Oct/18	
Status:	Treated
Project:	SDP Risk Register
Component/s:	C.1.2.2 Pipelines

Type:	Risk		
Reporter:	[6]	**Assignee:**	[8]
Labels:	Construction, PIP, risks		

Issue Links:				
	Relates			
	relates to	~~TSK-102~~	Calibration solver and MSMFS in dataflow	Closed
	relates to	~~TSK-2395~~	Compute Node Co-design: SDP sizing	Closed
Impact:	4 Major			
Probability:	4 Highly likely			
Exposure:	High			
Treatment:	Mitigate			
Treatment plan:	Continue benchmarking during construction.			
Residual Impact:	3 Moderate			
Residual Probability:	4 Highly likely			
Residual Exposure:	High			
Risk Event:	The implementation of processing components on hardware selected for SKA deployment has a lower than desired efficiency. Therefore certain processing components require more resources (e.g. memory bandwidth) than currently estimated leading to slower processing and therefore changes in the science observation schedule. The specific impact on the observation schedule will depend on the specific processing component that has a lower efficiency.			
Risk Cause:	Pre-CDR Prototyping is limited to components on current hardware. The residual risk remains that the complexity of implementing high performance processing components on emerging hardware leads to either a cost or schedule overrun, or within the SAFe development cycle a lower performance than desired is accepted.			
Existing Mitigation:	adopting G-GPU enabled modules from LOFAR and MeerKAT; ongoing work in CASA			

Comments

Comment by [9] (Inactive) [13/Apr/16]
Hello. This ticket has passed its planned end date and is still open. Please can you update the ticket (put planned end date in future or change the status) or let me know why it hasn't been resolved. Thanks.
Comment by [3] (Inactive) [22/Apr/16]
Who can we assign this to? [10]? Assign to [11] first and have her re-assign.
Comment by [11] [13/Sep/16]
Which version of the prototyping plan? Where is the relevant prototyping plan? I could not find this section in the one dated from 2015 Feb.
Comment by [2] [28/Sep/16]
Mitigation no longer relevant. We need to do some analysis (looking at the PT work?) to see how this happens.
Comment by [2] [11/Oct/16]
Work that is taking place includes:MSMFS and CLEAN algorithm parallelisation work ~~TSK-102~~Reports on reuse of existing code: TSK-85, ~~TSK-86~~, ~~TSK-87~~
Comment by [2] [24/Nov/16]
Refers to Risk SDPR-0067 in SDP RisksMitigation: As defined in the Adapting Software for Parallel Environments section of the Prototype Plan, this risk is to be fully explored in stage 2.Please define start and end date by 14 Aug 2015.Reassign the ticket if I made the wrong assignment.
Comment by [2] [24/Nov/16]
Time score: 4.0
Comment by [2] [24/Nov/16]
Fin score: 4.0
Comment by [2] [24/Nov/16]
Safety score: 1.0
Comment by [2] [24/Nov/16]
67.0
Comment by [2] [24/Nov/16]
Reputation score: 3.0
Comment by [2] [24/Nov/16]
performance score: 3.0
Comment by [2] [28/Feb/17]
title: pipeline modules are not able to exploit underlying multicore/SIMD/GGPU architectures. Risk: the need for individual pipeline modules to operate across nodes (not within a shared memory space) has not yet been defined, but may be required for some execution frameworks. Mitigation: adopting G-GPU enabled modules from LOFAR and MeerKAT; ongoing work in CASA. Change assignee to Ronald as Pipelines expert.
Comment by [2] [23/Oct/17]
The Calibration solver is a component that may require MPI or similar technologies and will need to communicate across multiple islands. Requires global synchronisation and does not map well to a data

flow architecture.

Comment by [2] [08/Jan/18]

We have mitigations in our heads, but we haven't captured them properly. We may have different phases in our workflow, and then we can optimise each phase individually. Also, formulate a new risk if appropriate.
We have tried to mitigate by our industry interactions. However, that has given rather inconclusive results. This is on-node performance - not between-node performance - which may need a separate ticket. Is this stalled because we're not particularly worried about this any longer?
Ferdl would like to see a list of pipelines modules, and whether they can or can't use this architecture efficiently?

Comment by [2] [21/Mar/18]

No new information. This remains the same.

Comment by [2] [26/Apr/18]

We are addressing this via the Vertical Prototyping work, which we will continue this sprint. However, we cannot yet rescore.
We were concerned with how we note the performance of various pipeline components. This will be addressed architecturally, not as part of this risk.

Comment by [8] [26/Apr/18]

A quick response to this:
So far vertical prototyping has looked at w-projection gridding and predict.
*We find that both fit well on GPU architectures and parallelism/concurrency can be exposed in different ways. ***
Independently John Romain and Bram Veenboer have produced some excellent work on IDG (See Bram's great talk here:
http://on-demand.gputechconf.com/gtc/2018/presentation/s8128-image-domain-gridding-on-accelerators.pdf)
All of this work demonstrates that these modules fit well on GPUs.

Further work that might be needed:

1. *Calibration is noted below, we could investigate this, but we have very little FTE effort (~0.8FTE at the moment).*
2. *I am not aware of work that optimises some of these algorithms (John and Bram have produced this for IDG) for multi-core/SIMD CPUs (such as Skylake). This might be worth investigating, but given performance on GPUs I'd suggest that it is a low priority).*

In summary I believe that the risk of finding an algorithm that is inherently sequential and has little or no parallelism that can be exploited by current architectures is low. __
Moreover it demonstrates that trying to quantify and algorithms performance in terms of achieved percentage of peak FLOPs is not a good metric for measuring performance or efficiency **

A slightly more detailed response:
Summary of vertical prototyping activities
Currently vertical prototyping has looked at "on-node" performance only. Specifically we are trying to define/understand:

1. Is the 10% efficiency assumption correct?
2. What are the limiting factors for any given algorithm (pipeline module). For example is an algorithm compute or bandwidth bound? Can it be broken down to fit into fast caches?...

3. What hardware considerations do we need to make for a given algorithm (pipeline module). For example how much memory does it need? If on an accelerator does it need a fast connection to the host (node).
4. Which (computational) algorithm (pipeline module) is most suited to any given accelerator (GPU/FPGA/CPU)?
5. Cost/Benefit of using native languages like CUDA compared to using pragma based approaches. For example do the performance benefits of a code that is specialised and more difficult to understand outweigh having a cleaner more maintainable code?

So far vertical prototyping has looked at some of these questions for w-projection gridding and predict. We intend to include a comparison of IDG (in collaboration with John and Bram).

Summary of vertical prototyping results
All of our results have been captured (thanks to Chris) in reports which can be found here: https://confluence.ska-sdp.org/display/WBS/Vertical+prototyping

Some important points to note so far:
For w-projection we compare two codes, one is compute limited, the other bandwidth limited (at L2 <-> L1). NOTE: the compute limited code is slower than the bandwidth limited code.
Take home message: Achieved percentage of peak FLOPs is not a good metric for measuring performance or efficiency of modern compute hardware.

Studying Brams excellent presentation on IDG: *http://on-demand.gputechconf.com/gtc/2018/presentation/s8128-image-domain-gridding-on-accelerators.pdf Slide 15 shows that for IDG to currently (this might change in the future) outperform w-projection a GPU using NVLink is required.*
Take home message: The choice of algorithms used in SDP pipelines greatly influences hardware requirements and therefore vertical prototyping such as the above is needed to help hardware architects deliver the correct architecture for SDP. This type of information can only be gained after an algorithm has been ported to an accelerator.

Wes.

Comment by [12] [27/Apr/18]

To add to Wes's comment, further details about the limiting factor for performance for the two w-projection codes have been added to the original report comparing these codes, in the performance section. The report is available under the original ticket TSK-2166 along with the full nvvp profiling reports for the two codes.

Comment by [2] [29/Jun/18]

Unchanged as of June 2018 RRB - work in this area has not yet concluded.
Suggest rewording this risk: title suggests that it's not possible to use such architectures at all, which isn't the case. Also risk event suggests that this is a design risk, rather than a construction/implementation risk. But the question is: do they become the processing-limiting step? It's very difficult to quantify now, as we can only benchmark against existing implementations of GGPUs. However, it's not clear that the scale of our problem at implementation will work with the new GGPUs. It can only be retired at construction.
This can only be mitigated by working with the vendors, and this doesn't completely alleviate this. We should note that the performance milestones in the Construction plan are to monitor and deal with this

risk. Also 10% efficiency here is misleading - it needs to be considered in the overall system context.
Comment by [2] [29/Jun/18]
Rewording of the risk - [6] and [13] to work on this.
Comment by [6] [12/Sep/18]
[13] to work on this. Perhaps needs new title. Event & cause needs to be updated. Mitigation linked to the 'Process/Plan to manage hardware technology risk' document.
Comment by [7] [19/Oct/18]
The residual impact is moderate because it does not affect the existence of the system but there is a moderate impact on the potential scientific throughput.

[SDPRISK-324] Buffers in the high performance ethernet network are overloaded Created: 20/Sep/16 Updated: 26/Sep/18	
Status:	Analyzed
Project:	SDP Risk Register
Component/s:	C.1.1.2.3 Bulk Data Network

Type:	Risk		
Reporter:	[27]	**Assignee:**	[14]
Labels:	Pre-Con, risks		

Issue Links:				
	Contains(WBSGantt)			
	contains	TSK-1746	Prototype data loss and buffer perfor...	Open
	Relates			
	relates to	TSK-872	Compute System: Requirements specific...	Open
	relates to	TSK-876	Compute System: Risk analysis	Open
Impact:	3 Moderate			
Probability:	2 Low Likelihood			
Exposure:	Medium			
Treatment plan:	Verify that tuning the network buffers at switch level will resolve this issue.			
Residual Exposure:	Medium			
Mitigation Start Date:	06/Jan/16			
Risk Event:	Data loss the SDP receive pipelines. System Performance/Science Impacted.			
Risk Cause:	The buffers in the bulk data network switches may be overloaded by traffic patterns common in the SDP.			

Comments

Comment by [15] (Inactive) [24/Nov/16]
SDP MT86 - 24May2016 (CB)
Comment by [2] [08/Dec/16]

Conventional switches, in particular cheaper models, may have limited and/or shared input buffers on the ports that receive data from CSP. Since the CSP data uses an unreliable protocol, overflow of such a buffer will cause irretrievable loss of data.In LOFAR we've observed that the buffers in current day switches may be overloaded by traffic patterns common in SDP. This may lead to data loss.
Comment by [2] [08/Dec/16]
110.0
Comment by [4] (Inactive) [23/Nov/17]
Hi [14], do you have suggestions for the next Risk Review Board meeting on this matter?
Comment by [14] [27/Nov/17]
This is still a risk that needs to be explored in more detail. Blocks on the availability of suitable hardware to test. Not high priority at the moment, since we're far from deployment.
Comment by [2] [21/Mar/18]
No change.
Comment by [7] [26/Sep/18]
[14] please could you provide an update on this risk by 5th October 2018? Many thanks.
Comment by [14] [26/Sep/18]
No change really. This is something to keep in mind when designing the network we're going to deploy and should probably be included in the acceptance tests for that system. This is still a risk, but not one we can fully retire so far from deployment (retire using a current device is useless, unless we deploy that exact device in production).

[SDPRISK-336] Receive Pipeline real-time Performance is Unachievable Created: 20/Sep/16 Updated: 19/Oct/18	
Status:	Analyzed
Project:	SDP Risk Register
Component/s:	C.1.2.1 Receive Components

Type:	Risk		
Reporter:	[27]	**Assignee:**	[14]
Labels:	Pre-Con, risks		

Issue Links:				
	Relates			
	relates to	TSK-775	Receive Components: Performance model...	Open
Impact:	3 Moderate			
Probability:	2 Low Likelihood			
Exposure:	Medium			
Residual Exposure:	Medium			
Risk Event:	Receive processing requires more resources than currently estimated leading to slower processing and therefore changes in the science observation schedule. No impact to cost or schedule.			
Risk Cause:	The Input Pipeline real-time Performance requirement is unachievable. (Computational)			

Comments

Comment by [15] (Inactive) [24/Nov/16]
"Converted OTA specific risk to general risk.Need to see if there is a similar risk here in the Pre-Con part of the risk Register. Also how do we track all our unknowns? for each WBS element and the associated interfaces?"
Comment by [2] [08/Dec/16]
The Input Pipeline real-time Performance requirement is unachievable. (Computational)
Comment by [2] [08/Dec/16]
115.0
Comment by [2] [21/Mar/18]

SIP have been working on this - review after 2018B sprint planning.
Comment by [2] [29/Jun/18]
Add mitigation to this ticket.
Comment by [7] [26/Sep/18]
[16] please could you comment on this risk by 5th October? Do you think there is an issue in this area given SIP findings or is it not possible to tell yet (in which case it remains open)? (Appreciate this is not well defined)
Comment by [16] [01/Oct/18]
Not possible to tell yet from prototyping. It was outside the remit of SIP to focus on assessing pipeline performance, so I can't really add anything useful to this risk with my SIP hat on. As you implied in your question, this risk is way too vague to be resolved in any useful way at this point given there will be multiple real-time pipelines within SDP with different properties and requirements. Clearly not all of those will be unachievable, and hopefully none of them will be so, but without requirements on individual pipelines and a lot more theoretical analysis or prototyping this is impossible to answer fully.

[SDPRISK-339] Data loss in Receive components Created: 22/Sep/16 Updated: 19/Oct/18	
Status:	Analyzed
Project:	SDP Risk Register
Component/s:	C1 Data Processor

Type:	Risk		
Reporter:	[27]	**Assignee:**	[14]
Labels:	Pre-Con, risks		

Issue Links:				
	Relates			
	relates to	TSK-780	Receive Components: Design verification	Open
Impact:	2 Minor			
Probability:	3 Likely			
Exposure:	Medium			
Treatment plan:	Over dimensioning part of the system should mitigate this risk.			
Residual Exposure:	Medium			
Mitigation Start Date:	06/Jan/16			
Risk Event:	Data loss in the SDP receive visibilities pipeline. Reduced Signal to Noise results in System Performance/Science Impacted.			
Risk Cause:	Based on experience of LOFAR, Data Processor could fail to meet the near real-time performance requirements (what are these?) and loses data.			
RRB Review comments:	Clarify which near real-time performance requirements are an issue. Give more detail on what happens when the issue occurs (event). 15/09/2016 FG & CB			
Existing Mitigation:	Over dimensioning part of the system it should mitigate this risk.			

Comments

Comment by [15] (Inactive) [24/Nov/16]
SDP Data Processor Design Doc (CB)
Comment by [2] [08/Dec/16]
Based on experience of LOFAR, Data Processor could fail to meet the near real-time performance requirements and loses data.What are the "near real-time performance requirements"? Are they

defined? Or is this risk On-Hold until the Risk "near real-time
Comment by [2] [08/Dec/16]
Risk Number: 106.0
Comment by [2] [28/Feb/17]
title: data loss in Receive components
Comment by [2] [21/Mar/18]
Check with SIP at sprint planning, may be able to rescore.
Comment by [2] [29/Jun/18]
add mitigation to ticket.
Comment by [7] [26/Sep/18]
[16] please could you comment on this ticket by 5th October based on SIP prototyping work/conclusions? Thank you.
Comment by [16] [02/Oct/18]
Given this risk is fairly poorly defined as depends on the fraction of per node BDN bandwidth required, and what the per node BDN bandwidth is specified to be, it's hard to give an assessment which holds much weight from the analysis done in prototyping. Having said that, we have had good success in prototyping the visibility receive interface in SIP (which will be described in the SIP memo), but this required a lot of work and careful configuration of P3 to achieve. Prototyping has also highlighted some outstanding issues in the best way to use the current version of the SPEAD library to be able to maximise UDP line rates on the P3 25 Gb/s BND - do this requires very large numbers (100's) of UDP streams which have to be handled concurrently which is difficult and/or resource intensive. I should also note (thanks to Fred for pointing this out to me) that TCP has very recently been added to the SPEAD library. While we have not yet had a chance to test this prototcol with the SIP receive code, the use of TCP with SPEAD is described in the excellent SDP Memo 046 (https://confluence.ska-sdp.org/download/attachments/201294049/SDP%20Memo%20046_%20Experiences%20with%20the%20SPEAD%20protocol.pdf?version=1&modificationDate=1527600998000&api=v2). As noted in this memo, use of TCP may help reduce some aspects of the risk, but the ICD currently specifies that the prototcol used will be UDP, so I think there is proably more analysis and prototyping to be done. Based on prototyping experience, I would therefore suggest to either leave this as a medium risk, or raise the risk slightly. The prototyping experience has shown that while the problem of managing data loss is well defined, it will very likely take some considerable development effort in the production SDP to make to make this work and solve all remaining implementation issues.
Comment by [14] [03/Oct/18]
Note that moving to the TCP/IP version of SPEAD for the CSP-SDP links would require significant more resources (both FPGA real estate and programming) on the CSP side to implement. TCP/IP requires handshaking and bi-directional communication, whereas UDP/IP does not.
Comment by [7] [19/Oct/18]
RRB27: Reviewed recent comments. We believe the risk scoring should be left unchanged.

[SDPRISK-343] Incomplete Interface Description between pipelines components & Execution Framework Created: 20/Sep/16 Updated: 19/Oct/18	
Status:	Treated
Project:	SDP Risk Register
Component/s:	C.1.2.2 Pipelines, C.1.2.3 Execution Framework

Type:	Risk		
Reporter:	[27]	**Assignee:**	[6]
Labels:	Interface, Pre-Con, SE, risks		

Issue Links:				
	Relates			
	relates to	~~TSK-52~~	SE: L4 interface definition, ICDs, etc.	Resolved
	relates to	TSK-330	High-level architecture documentation	In Progress
	relates to	TSK-1288	Sprint data model tasks	In Progress
	relates to	~~TSK-1840~~	Using ARL components with DALiuGE EF	Closed
	relates to	~~TSK-1830~~	Develop generalised interfaces to ARL...	Resolved
	relates to	~~TSK-1754~~	Execution Frameworks Pipeline compone...	Closed
	relates to	~~TSK-1829~~	Map critical algorithms into Dask	Closed
Impact:	4 Major			
Probability:	5 Near Certainty			
Exposure:	Extreme			
Treatment:	Mitigate			
Treatment plan:	Although generic solutions have been identified further detailed design and prototyping is required to assess the cost and complexity of using these solutions.			
Residual Impact:	3 Moderate			
Residual	3 Likely			

Probability:	
Residual Exposure:	Medium
Mitigation Start Date:	06/Jan/16
Risk Event:	Potential Schedule overrun or using cost contingency due to additional effort needed in the development of the Processing Components or the Execution Frameworks.
Risk Cause:	Interface between processing components and Execution Framework not clearly enough specified. We do not have a generic interface between pipeline components and the Execution framework (although we have wrapped some pipelines components for use by DaLiuGE).
RRB Review comments:	Ferdl to reword risk cause. 28-9
Existing Mitigation:	Existing pipeline components have been successfully interfaced to the DFMS, albeit via wrapping some components.

Comments

Comment by [6] [22/Sep/16]
This interface has not been specified yet, so this looks like a task that should be on the WBS rather than a risk or failure mode.
Comment by [2] [28/Sep/16]
This should be added as a task to the WBS.
Comment by [2] [28/Sep/16]
This needs the risk cause rephrasing, as this is a risky interface.
Comment by [2] [08/Dec/16]
Interface between PIP components and EF not clearly enough specified.
Comment by [2] [08/Dec/16]
Time score: 5.0
Comment by [2] [08/Dec/16]
Fin score: 3.0
Comment by [2] [08/Dec/16]
Safety score: 1.0
Comment by [2] [08/Dec/16]
Risk Number: 104.0
Comment by [2] [08/Dec/16]
Reputation score: 2.0
Comment by [2] [08/Dec/16]
performance score: 5.0
Comment by [2] [28/Feb/17]

expand abbreviations in title. pluralise component. Mitigation of interface to DFMS isn't correct. DFMS has wrapped some components, but the risk here is that we need to define a generic interface between pipeline components and the execution framework, and this has not been done. Treatment: design this interface. Consideration needs to be given to baseline execution framework and other potential execution frameworks. To derisk, we need a generic definition. Probability: Near Certainty (because we have not done the work). Impact: major (time delay)
Comment by [2] [02/Jun/17]
Work is proceeding on this, but is not yet complete. This should be a priority for 2017C.
Comment by [2] [31/Aug/17]
This should also receive attention in this current sprint from the architecture (to write up) and prototyping teams to test how we can interface to the ARL from a number of languages, plus accessing CASA components. There will be a number of tickets in the next sprint.
Comment by [7] [01/Sep/17]
We have reduced risk in 2017C (with introduction of wrappers of components that have an architectural description of the interface) but will reduce much further when wrapper approach fully demonstrated by prototyping in 2017D.
Comment by [4] (Inactive) [22/Nov/17]
"We haven't seen much result on this (despite having said previously that this was a priority for 2017C) - we can only rescore after prototyping. While we have a design for this, **we haven't verified by prototyping that we've resolved the issue**" quote from RRB21 23OCT17
Comment by [4] (Inactive) [03/Jan/18]
Being close to the next Risk Review Board meeting, I wonder if you have anything to state or to share that could facilitate the progress on the risk management or could better describe the status quo to the watchers. By the way, Happy New Year!
Comment by [2] [08/Jan/18]
Work is ongoing. The architecture work on this remains incomplete and will carry on into the next sprint, though we've made progress. The relevant work has been on the pipelines document, but we need to take that work and analyse it carefully, and ensure that it's consistent through the architecture. Do we understand the details of the data models well enough to derisk everything at this point? We have in-memory data models, which we hope will give adequate performance. However, we haven't defined them. This is fairly reliant on the ARL. We have many related issues, but no one piece of work will solve this. Need a story to take forward the architectural work here, with a vertical deep dive. Analyse the contributing work, and draw out a consistent architecture. (note, this is separate from the MSv3 work)
Comment by [2] [21/Mar/18]
A lot of work went in over the last sprint. Return to this at the end of the PCG, once we've reviewed our progress.
Comment by [2] [26/Apr/18]
This will drive a lot of work this sprint. We are on course to reduce risk by CDR. However, we need to verify the architecture we have by prototyping.

Comment by [2] [29/Jun/18]
Architecturally, we have addressed this risk. This aspect is complete, and documented in Rev 05 of the Architecture. The June 2018 ATAM produced few comments on this area. The technical difficulties with this interface, which interfaces with the JVM and other systems, have been prototyped. While the prototyping is progressing, issues remain. However, we see a route to an implementation that addresses the problems in this area. [5] would be able to provide a more detailed update. It's still highly likely that there will be issues with the implementation. However, the impact should be lower. We've made progress on managing the metadata (which was one of our driving concerns), and we think that there is a route to interfacing to the components (though it may not be pretty). Impact now moderate. Work will continue in the next sprint.
Comment by [2] [04/Sep/18]
Comment from today's EF meeting: Can't reduce the risk - the interface is pretty horrible to work with, and not complete. The mitigation plan may not succeed in further reducing the risk. We don't think it's possible to rescore the risk.
Comment by [6] [12/Sep/18]
Task in the next sprint: Analyse the issue to determine what the new mitigation plan needs to be. Also see if this might impact the architecture or only this particular interface.
Comment by [7] [19/Oct/18]
Progress made on wrapping C components (that can be called from any framework). The route has been identified. The route to managing data across the interface has also been identified (protocol buffers). Detailed design is still required but we can see ways to do it so there is still some significant residual risk.

[SDPRISK-344] Data rate between nodes is underestimated Created: 09/Sep/16 Updated: 19/Oct/18	
Status:	Treated
Project:	SDP Risk Register
Component/s:	C.1.1.1.1 Compute Island, C.1.2.2 Pipelines

Type:	Risk		
Reporter:	[27]	**Assignee:**	[5]
Labels:	PROT, Pre-Con, risks		

Issue Links:				
	Relates			
	relates to	TSK-867	Compute System: Performance model	Open
	relates to	TSK-370	Performance model	In Progress
	relates to	TSK-1303	P3 SIP work	In Progress
Impact:	4 Major			
Probability:	3 Likely			
Exposure:	High			
Treatment:	Mitigate			
Treatment plan:	Modelling data transfer rates in the Performance Model and further analysis of algorithmic choices for workflows that may need high data rates between nodes. Added Performance Verification milestone in the SDP Construction Plan.			
Residual Impact:	4 Major			
Residual Probability:	2 Low Likelihood			
Residual Exposure:	Medium			
Mitigation Start Date:	02/Feb/15			
Risk Event:	Certain workflows require higher data rates between nodes than currently estimated causing these workflows to perform slower. This will cause changes in the science observation schedule. The specific impact on the observation schedule will depend on the specific workflow. No impact to cost or schedule.			
Risk Cause:	The required data transfer rate between nodes and data islands has been underestimated in the baseline hardware design.			
RRB Review comments:	mitigation needs to be updated (after risk analysis) 28-9 FG.			

Comments

Comment by [15] (Inactive) [24/Nov/16]
Reasses based on WIP
Comment by [2] [08/Dec/16]
The required data transfer rate between nodes and data islands has been underestimated in the baseline hardware design.
Comment by [2] [08/Dec/16]
Time score: 3.0
Comment by [2] [08/Dec/16]
Fin score: 4.0
Comment by [2] [08/Dec/16]
Safety score: 1.0
Comment by [2] [08/Dec/16]
23.0
Comment by [2] [08/Dec/16]
Reputation score: 2.0
Comment by [2] [08/Dec/16]
performance score: 3.0
Comment by [2] [28/Feb/17]
Title: Data rate between nodes is underestimated
Comment by [2] [23/Oct/17]
Review this risk in the light of developments in our design and prototyping.
Comment by [2] [26/Oct/17]
Can we retire this risk, given we've costed a fully bisectional bandwidth? Please provide a comment by 01 November.
Comment by [7] [02/Nov/17]
John's view is that there are still algorithmic choices to make to fully exploit the topology of the network. PWo has a model giving data rates for various cases. The inter-node data rate is dependent on the algorithm choice and therefore the risk can not be retired but may be rescored.
Comment by [4] (Inactive) [03/Jan/18]
Being close to the next Risk Review Board meeting, I wonder if you have anything to state or to share that could facilitate the progress on the risk management or could better describe the status quo to the watchers. Thank you.
Comment by [2] [08/Jan/18]
Impact could still be major - if we've got the data rate wrong, our entire architecture could underperform. Whilst we're unsure of what the data rate is, the impact will still be major. We've improved our understanding by analysis of what the data rate is likely to be, so we're more confident that it's not likely to happen. Where we have found issues (e.g. with the calibration solver), we will address those in a different way. If we do have to do a massive data corner turn, it could still be

a big impact on performance, but we are tackling this in other risks, and are managing it there. Therefore, we can reduce the probability to Low likelihood (2).
Comment by [2] [21/Mar/18]
No further information.
Comment by [7] [19/Oct/18]
This will remain into the detailed implementation of workflows, with realistic data rates and platform scale, into construction.

[SDPRISK-346] Degree of required spatial coverage for Fast Imaging pipeline is uncertain Created: 12/Sep/16 Updated: 19/Oct/18	
Status:	Analyzed
Project:	SDP Risk Register
Component/s:	C.1.2.2 Pipelines

Type:	Risk		
Reporter:	[27]	**Assignee:**	[1]
Labels:	PIP, Pre-Con, risks		

Attachments:	sdp-transients-pipeline.pdf
Issue Links:	**Relates** relates to TSK-1144 Pipelines: Imaging Analysis Open
Impact:	3 Moderate
Probability:	3 Likely
Exposure:	Medium
Treatment plan:	Pursue clearer requirements and expectations from the SKAO.
Residual Exposure:	Medium
Mitigation Start Date:	01/Dec/15
Risk Event:	The Fast Imaging workflow will require more resources (compute) than currently estimated leading to changes in the science observation schedule. No impact to cost or schedule.
Risk Cause:	Fast Imaging compute load uncertainty and applicability of SKA1-SYS_REQ-2345: Uncertainties remain in our model of the Fast Imaging pipeline and in the degree to which this requirement is meant to be implemented at full spatial and temporal resolution.

Comments

Comment by [6] [14/Sep/16]
[3] to review the status of this risk in light of the memo by [17] from Feb 2015 and d-PDR design document. Update mitigation as well (with further work if required).
Comment by [2] [08/Dec/16]

Fast Imaging compute load uncertainty and applicability of SKA1-SYS_REQ-2345: Uncertainties remain in our model of the Fast Imaging pipeline and in the degree to which this requirement is meant to be implemented at full spatial and temporal resolution.To Do: [3] to review the status of this risk in light of the memo by [17] from Feb 2015 and d-PDR design document. Update mitigation as well (with further work if required).
Comment by [2] [08/Dec/16]
time score: 1.0
Comment by [2] [08/Dec/16]
Fin score: 3.0
Comment by [2] [08/Dec/16]
Safety score: 1.0
Comment by [2] [08/Dec/16]
34.0
Comment by [2] [08/Dec/16]
Reputation score: 2.0
Comment by [2] [08/Dec/16]
performance score: 3.0
Comment by [3] (Inactive) [11/Jan/17]
This will be progressed under ~~TSK-75~~
Comment by [2] [28/Feb/17]
title: degree of required spatial coverage for fast imaging pipeline uncertain.
Comment by [2] [08/Jan/18]
Link to the other Fast Imaging risk - they both relate to uncertainty in how the Fast Imaging science requirements are defined.
Comment by [6] [26/Sep/18]
[1], please look at this risk together with SDPRISK-320.
Comment by [6] [26/Sep/18]
Review again at next RRB.
Comment by [7] [19/Oct/18]
RRB27: No change to be recorded as per risk SDPRISK-320. Further work in the bridging period recommended.

[SDPRISK-359] Processing/Calibration Strategy is not well defined Created: 28/Jul/16 Updated: 19/Oct/18	
Status:	Treated
Project:	SDP Risk Register
Component/s:	C.1.2 Processor Software

Type:	Risk		
Reporter:	[2]	**Assignee:**	[5]
Labels:	Architecture, Pre-Con, risks		

Attachments:	L1 rev9 new requirements.xlsx			
Issue Links:				
	Relates			
	relates to	~~TSK-32~~	Calibration Strategy	Closed
	relates to	TSK-1243	Pipelines: Calibration strategies	In Progress
	relates to	~~TSK-102~~	Calibration solver and MSMFS in dataflow	Closed
	relates to	TSK-1849	Organise meeting to discuss calibrati...	Closed
	relates to	~~TSK-77~~	Provide report on current state of SK...	Resolved
	relates to	~~TSK-75~~	SDP contribution to Calibration strat...	Closed
	relates to	~~TSK-76~~	provide resource for real-time calibr...	Closed
	relates to	SDPRISK-340	Inter-compute island communications i...	Closed
Impact:	5 Catastrophic			
Probability:	5 Near Certainty			
Exposure:	Extreme			
Treatment:	Mitigate			

Treatment plan:	Mitigation: Make a reasonable assumption for the calibration strategy until the final calibration strategy is defined at system level (based on Tim's work). Work with SKAO to define the calibration strategy at system level. Also develop the architecture so detailed knowledge of the calibration strategy is not needed to support any current calibration strategy used on any facility.
Residual Impact:	3 Moderate
Residual Probability:	2 Low Likelihood
Residual Exposure:	Medium
Mitigation Start Date:	06/Jan/16
Risk Event:	Insufficient compute resource (since performance was underestimated) for certain calibration observations will lead to reduced scientific performance/quality or reduced scientific output due to changes in the science schedule.
Risk Cause:	Processing/Calibration Strategy is not well defined causing the computational performance to be underestimated. Final Calibration update rates only known when hardware is in the field.
RRB Review comments:	Covered partially by SDPR-M13-G24. Updated cause & score. 26/7 FG

Comments

Comment by [6] [26/Sep/16]
The L1 rev9 requirements release does include a few new requirements relating to calibration: L1 rev9 new requirements.xlsx. I've not been able to find the references or parent documents for these requirements. I'm requesting further info and reference material from Juande.
Comment by [6] [26/Sep/16]
[5], based on the discussion we had today, we need to update this risk and review the scoring. I will link any new calibration related material to this issue as soon as I become aware of it.
Comment by [5] [26/Sep/16]
Somebody should analyse the new requirements ASAP to see what the consequent L2s might be but clearly they will not address all the questions. We should increase the scoring to performance/cost 5 and delay 5. The reasons are as follows:* Calibration has a big impact on both FLOP requirements and interconnect and software performance requirements. For example our power budget is 70% for DFT for calibration; and the most challenging step for graph based approaches is the calibration step which is still poorly defined.* Calibration is complex to analyse, parallelise and make efficient. Incorrect design will undoubtedly contribute to large (multi year) delays. Calibration strategy is a critical dependency of many of the highest risk items in our work, for example:* One of the most challenging processing steps for the execution engine design is the calibration step; this however is not well defined, specifically what is the smoothness across frequency of the calibration solutions and/or at what time scale and across what bands effects such as TEC and delay should be fit* The information that is stored in the LTS, and the volume of intermediate data products related to calibration is dependent on the calibration

strategy. Both of these are big current risks. * We currently do not have a parallelisation strategy for the calibration solver step and this depends on the calibration strategy* The interconnect requirements between the compute island might be dominated by calibration solution communication and/or full bandwidth synthesis for calibration purposes.
Comment by [2] [28/Sep/16]
Risk score updated, now has score of 25, based on the scoring Bojan proposed. Likelihood is 5 - they cannot produce an adequate calibration strategy in the next 6 weeks, before System PDR is due. This is blocking our progess on our L2 requirements.
Comment by [2] [04/Nov/16]
Bojan's previous analysis (23/06/2016) suggested moving this to the Construction Risks, with a revised scoring. (Revised score not specified)
Comment by [5] [04/Nov/16]
Did I understand correctly that we were going to abandon the distinction between construction and pre-construction risks?
Comment by [2] [04/Nov/16]
We're not abandoning it; however, there aren't that many technical risks that are realised in pre-construction. (i.e. the risk is we deliver a poor design, but we don't discover this until we try to build it.) We can't abandon the distinction altogether, because management risks **do** fall neatly into pre-construction and Construction categories.
Comment by [15] (Inactive) [24/Nov/16]
"PRO-32 On-Hold: To Be Abandoned?"
Comment by [2] [24/Nov/16]
Processing/Calibration Strategy is not well defined at system PDR (and therefore does not leave SDP enough time to take it into account in the design).Comments: Covered partially by SDPR-M13-G24. Updated cause & score. 26/7 FGComment: 0095 move to technical preconstruction risks. Cause updated: 'therefore causing uncertainty in the design (e.g number of major cycles) - could lead to reduced scientific performance. Noting that poor definition at System PDR carries risk of SDP not being able to take calibration design into consideration. likelihood 5: performance cost 4, time delay 4. Score: 16 (high)
Comment by [2] [24/Nov/16]
Time score: 5.0
Comment by [2] [24/Nov/16]
Fin score: 3.0
Comment by [2] [24/Nov/16]
Safety score: 1.0
Comment by [2] [24/Nov/16]
95.0
Comment by [2] [24/Nov/16]
Reputation score: 2.0
Comment by [2] [24/Nov/16]
performance score: 5.0
Comment by [2] [28/Feb/17]

alter title : remove "at system pdr" otherwise this remains as is.
Alter labels - remove PIP, add architecture.
Proposed mitigation: work with SKAO on a resolution team to resolve this. External constraints mean this hasn't happened yet, because of other priorities. We've also done preliminary work on distributed calibration.
Link to calibration epic.

Comment by [2] [02/Jun/17]

There has been some work on this. However, it's not mature enough yet to review the scoring.

Comment by [7] [31/Aug/17]

RRB comment: This situation behind this risk will be reviewed during 2017D. This together with work in progress will impact the risk rating and help derive more specific risks.

Comment by [7] [01/Sep/17]

PCG: We have evolved the SDP architecture to provide flexibility. As a outcome of a SDP system wide quality attribute workshop we formalised flexibility for workflow construction and deployment. This has been built into the top-level architecture in particular the module structure. As a result the lack definition now should only result in requirements to design additional processing components for the new work flows. The cost of these additions it mitigated further by the SAFe process during construction.
Additionally analysis of key algorithm types in particular solvers has moved forward and the impact assessed on the architecture has been reduced - for example the architecture could accommodate a global MPI process.
The SDP-low calibration working group and other improvements in the L2 requirements have not identified problems which would require a change of architecture.
The SKAO now has a dedicated interferometry expert working on calibration as well.

Comment by [4] (Inactive) [22/Nov/17]

Quote from the RRB21 23OCT17 "*We haven't made further progress on this. There are still uncertainties, but it should decrease as we progress to CDR. We need a Calibration meeting with Robert TSK-1849*"
.

Comment by [4] (Inactive) [03/Jan/18]

Being close to the next Risk Review Board meeting, I wonder if you have anything to state or to share that could facilitate the progress on the risk management or could better describe the status quo to the watchers (have we made progress on this?). By the way, Happy New Year!

Comment by [2] [08/Jan/18]

Have still not had meeting with [24], as he's very busy.

Comment by [2] [21/Mar/18]

Are scheduling meetings with [24]. Awaiting outcome.

Comment by [2] [26/Apr/18]

Meeting with [24] planned, but ongoing work will be required. We may carry this risk through CDR, as there isn't the SKA resource to allow us to resolve. Nor is there much domain expertise in this risk.
There should be a task this sprint to see whether there is more that we can do to mitigate this risk via the architecture. We could supply SKA with information about what we can supply to them e.g. cadence of calibration solutions, to help with the calibration strategy.
This is very hard to reduce before we start taking data, because we are in a new regime for signal to noise. Therefore we review the risk as part of the sprint to see what is sensible for the SDP to do by CDR.

Comment by [2] [29/Jun/18]
Need to update risk cause and mitigation, but this can be done offline ([6])
Comment by [2] [29/Jun/18]
SDP will work on this with SKAO, but this risk will remain for some time, as it needs to be worked on in an iterative fashion.
Comment by [6] [04/Sep/18]
Updated the Risk Cause and Event.
Comment by [6] [04/Sep/18]
Removed text from **Existing Mitigation** field as it seems to overlap with the text in the **Treatment plan** field (which it more up-to-date).
Comment by [6] [12/Sep/18]
This issue was discussed during the pre-CDR review and afterwards. Progress has been made, but the output has not been documented yet. Need to follow up on who is documenting the this and when the document is expected. Once the document is available we can reassess the risk and re-score if necessary.
Comment by [7] [19/Oct/18]
We have made progress and can rescore. We have not defined a better strategy but abstracted all the potential ways of doing calibration in literature into two ways we have demonstrated can be supported by the architecture. So our confidence has increased. This is addressed in the architecture workflow views. In particular the model partition calibration view. Likelihood is much reduced. We do not know the computational cost but do know the architecture can support it. This is a tradeoff of scheduling. Impact on science will now be modest rather than major now.

[SDPRISK-360] Uncertainty or inaccuracy in the Parametric Model computational estimate Created: 28/Jul/16 Updated: 19/Oct/18	
Status:	Treated
Project:	SDP Risk Register
Component/s:	C Science Data Processor

Type:	Risk		
Reporter:	[2]	**Assignee:**	[6]
Labels:	PROT, Pre-Con, risks		

Attachments:	continuum_verification.pdf		
Issue Links:	**Relates**		
	relates to	TSK-370	Performance model In Progress
	relates to	~~TSK-223~~	SE: Cost submission Closed 10 Feb
Impact:	4 Major		
Probability:	5 Near Certainty		
Exposure:	Extreme		
Treatment:	Mitigate		
Residual Impact:	2 Minor		
Residual Probability:	5 Near Certainty		
Residual Exposure:	Medium		
Mitigation Start Date:	03/Feb/15		
Risk Event:	Certain workflows require more resources (compute & buffer) than currently estimated leading to changes in the science observation schedule. The specific impact on the observation schedule will depend on the specific workflow that was underestimated by the Parametric Model. No impact to cost or schedule.		
Risk Cause:	The verification of Parametric Model indicates that the parametric model underestimates the actual performance required due to the parametric model being incomplete or inaccurate.		
Existing Mitigation:	Continue updating the parametric model to incorporate analytic expressions for all of the key processing bottlenecks. SKAO RT to analyse these related issues and provide updated inputs and changes to the Parametric Model & System Sizing.		

Comments

Comment by [2] [04/Nov/16]
[5] suggested transferring this risk to construction and creating a new risk for pre-Construction of: SDP design precludes implementation of key kernels with computational efficiencies approaching the theoretical maximum and Insufficient down-selects in SDP design preclude implementation of efficient computational kernels in the construction phase
Comment by [5] [08/Nov/16]
This comment was about the old M8-P32 and M8-P18 risks but this ticket looks to me to be the old M8-24 risk.
Comment by [5] [08/Nov/16]
The title of this should be "Parametric Model underestimates the computational requirements" . The impact would be that we are able to support less science than planned. It is unlikely there would be a cost impact since we are unlikely to buy more hardware to compensate. I don't think moderate errors in the Parametric Model will change the design or architecture in a way – what was the rationale for the design impact?I don't think the SWE work for verification is continuing in any way, so this link needs to be revised. Furthermore ASKAPSoft does not deal with key aspects of our model, such as facetting and A-projection. So we haven't planned a mitigation at this time.I think the likelihood is indeed high. Not sure about the impact, because reasonable errors will not I think affect the design. What was rationalle for a high time impact?
Comment by [5] [08/Nov/16]
I've commented on this risk. It needs to be revised.
Comment by [15] (Inactive) [24/Nov/16]
"PRO-32 On-Hold: To Be Abandoned?"
Comment by [2] [24/Nov/16]
Fin score: 4.0
Comment by [2] [24/Nov/16]
Safety score: 1.0
Comment by [2] [24/Nov/16]
31.0
Comment by [2] [24/Nov/16]
Reputation score: 2.0
Comment by [2] [24/Nov/16]
performance score: 4.0
Comment by [15] (Inactive) [24/Jan/17]
[5] Attached memo to the risk. Please re-score as appropriate.
Comment by [5] [24/Jan/17]
I think we should re-score the likelihood one notch since some verification is done. Note however no current mitigation of this is ongoing, and the verification work is now 1.5 years out of date – our parametric model has moved on. So at the present time I don't think we are mitigating this risk. Additionally, at the moment we have an identified **issue** in that the parametric model does not cover peeling.*

Comment by [2] [02/Jun/17]
Unless pipelines are fully defined, the parametric model may underestimate the computational requirements. Until pipelines are defined, this risk remains. A possible avenue to reduce this risk could be to compare with existing software that does similar work to SDP at a smaller scale. However, it's not clear that this will give clear answers.
Comment by [2] [08/Jan/18]
Suggest more work on the parametric model on the current sprint.
Comment by [2] [21/Mar/18]
[28] has run some more experiments with the parametric model, with more observation scenarios. This hasn't exposed any weaknesses in the parametric model. An SKA document has been produced. However, there isn't enough evidence to rescore at this stage.
Comment by [2] [29/Jun/18]
Need to update the wording. Note that this is different to risk 401 - this risk is related to the input that we need from SKA (i.e the workflows we need to run etc). We need to do analysis first. We can't prototype in SDP every single pipeline. Need clarifications about how we deal with cost overruns. This will be worked on by the resolution team. This risk should not have been transferred, as it's about the accuracy of the parametric model, which is in our domain. We need to continue updating the parametric model to incorporate analytic expressions for all of the key processing bottlenecks.
Comment by [7] [29/Jun/18]
Here is the comment from the August 2017 RRB: "360 - this will remain.However, impact may be lower now that we've defined some of the procurement process. (the initial SDP procurement is a fixed cost item - therefore impact of getting sizing wrong is only impact on the science, it's not on cost or project schedule.) rescore - impact one notch down. Risk transferred from SDP onto Science Operations."
Comment by [6] [04/Sep/18]
Proposed new working for review by RRB: **Risk Event:** Certain workflows require more resources (compute & buffer) than currently estimated leading to changes in the science observation schedule. The specific impact on the observation schedule will depend on the specific workflow that was underestimated by the Parametric Model. No impact to cost or schedule. **Existing Mitigation:** Continue updating the parametric model to incorporate analytic expressions for all of the key processing bottlenecks. Remove the current text in the **Treatment plan** field. Its not clear why this risk has a probability of 5. The Risk Cause should include the rationale for the probability.
Comment by [6] [12/Sep/18]
[18] made the point that the impact of this risk is on the observation schedule and doesn't impact on any L1 requirements, cost or schedule (during construction) and therefore the Impact score of the risk may be lower than 3. Previous RRB recommended to reduce impact to 2.
Comment by [7] [19/Oct/18]
RRB27: The title of the risk was changed to better reflect the risk cause and event.

[SDPRISK-363] Buffer hardware and software does not meet performance requirements Created: 19/Jul/16 Updated: 19/Oct/18	
Status:	Treated
Project:	SDP Risk Register
Component/s:	C.1.1 Processor Platform

Type:	Risk		
Reporter:	[2]	**Assignee:**	[19]
Labels:	Architecture, EF, Pre-Con, risks		

Issue Links:				
	Relates			
	relates to	~~TSK-1551~~	Test standalone module on P3-AlaSKA	Closed
	relates to	TSK-1307	Buffer modelling	On hold
	relates to	TSK-1833	Buffer Hardware and Software Design	Open
	relates to	TSK-256	Define object storage requirements	Closed
	relates to	~~TSK-557~~	SIP: Update visibility data ingest	Closed
	relates to	~~TSK-1626~~	SIP: Develop Visibility Ingest module	Closed
Impact:	4 Major			
Probability:	3 Likely			
Exposure:	High			
Treatment:	Mitigate			
Treatment plan:	Contract on IO performance, P^3 system for prototyping.			
Residual Impact:	3 Moderate			
Residual Probability:	2 Low Likelihood			
Residual Exposure:	Medium			
Risk Event:	System Performance will be impacted. Architecture and design may need to be			

	changed leading to use of software cost contingency (schedule or cost). Buffer hardware may need to be changed leading to use of hardware contingency. Contingency already included in cost estimate.
Risk Cause:	I/O Interface insufficiently supports data parallelism at our cost point at 2024-2025
Existing Mitigation:	Continue prototyping until construction phase. We will continue to benchmark against available hardware. See hardware mitigation plan in cost basis of estimate.

Comments

Comment by [5] [22/Jul/16]
I think the action here is to clarify cause and event. I'm assigning to Andreas as it appears he is the author.
Comment by [20] [12/Sep/16]
I was not the author of the wording here. The original risk wording was different. Having said this, I'm now looking into mapping what we've already done in this regard to this risk and then analyse the remaining risk, perform additional prototyping and eventually propose a change of the score. The wording of this risk implies the extreme approach to the granularity of the tasks, which in turn moves all the I/O from within the task to the framework. I would suggest to add a risk regarding the task granularity, because this risk is just a consequence of the task granularity.
Comment by [5] [12/Sep/16]
I think it worth if can find the author so that we make sure we cover what they intended in this risk. [6] / [7] do you know who originally drafted this risk ?
Comment by [20] [12/Sep/16]
The original text is in the Execution Framework dPDR document, but it does not mention the data parallelism. Thus it now has a different scope and context. ----I/O Interface insufficient: We judge this risk as ‘very likely’ for the current DFMSimplementation, and ‘probable’ for the first real implementation. The architecture itself isvery flexible in this regard and allows to ‘plugin’a wide variety of existing or upcomingI/O frameworks. The amount of work required to identify, interface and optimise suchframeworks is a separate related, but lower level risk. We are working on a number offrameworks, such as ADIOS, HDF5 and some advanced database engines. Again thisrequires more work and prototyping, but is planned and can be distributed. The impact isvery high, since it will limit the amount of science that can be done.
Comment by [6] [14/Sep/16]
The risk cause needs to be expanded as "I/O Interface insufficient" does not fully describe what the problem is. Its possible that I added the part on data parallelism to try and expand on the risk cause as there was a very similar risk on the risk register. I've created an A3 page for this risk, please use this to provide more detail on this risk.https://confluence.ska-sdp.org/pages/viewpage.action?pageId=196444365
Comment by [2] [24/Nov/16]
I/O Interface insufficiently supports data parallelismScore: 25 (critical) 19/07/2016Assigning to Bojan as Project Engineer
Comment by [2] [24/Nov/16]
Time score: 4.0
Comment by [2] [24/Nov/16]

Fin score: 4.0
Comment by [2] [24/Nov/16]
Safety score: 1.0
Comment by [2] [24/Nov/16]
97.0
Comment by [2] [24/Nov/16]
Reputation score: 2.0
Comment by [2] [24/Nov/16]
performance score: 4.0
Comment by [2] [28/Feb/17]
Given our architecture and execution framework, can the buffer performance and IO performance be sufficient to support our chosen execution framework. Is this what this risk was about? Assign to John to consider further, and to help us define the risk wording. Treatment: Canonical Contract on IO performance. , P^3.
Comment by [2] [01/Mar/17]
John: please consider this risk further, and help us refine the wording. I'm happy to discuss with you.
Comment by [19] [01/Mar/17]
Does the "I/O interface" refer to ingest (write to buffer) or does this refer to both read and write from the buffer.
Comment by [2] [01/Mar/17]
I'd assume it refers to read and write - the I/O isn't one way. If there are distinct risks associated with reading and writing, then we can adjust this risk so it deals with one of these, and create a new risk for the other.
Comment by [2] [02/Jun/17]
e..g some buffers can solve this by striping. This is what we'd like to benchmark. Impact raised to catastrophic - this would mean our buffer couldn't work. Our entire architecture of how we're using the buffer to share data between nodes wouldn't work. because it's a software problem, buying more buffer wouldn't help. This is then a cost and schedule issue - it would be huge work to rearchitect the system, or we'd have to accept a high performance impact. Until we've got good benchmarks and so on, we can't reduce the likelihood.
Comment by [2] [02/Jun/17]
[5] please provide better description.
Comment by [19] [24/Aug/17]
I have started work to try and define some of the prototyping activity. I am capturing this work here, which is very much WIP https://confluence.ska-sdp.org/pages/viewpage.action?spaceKey=WBS&title=Batch+Processing From my understanding of some of the Views and Beyond architectural work the Buffer module now subsumes a number of storage requirements, not just the batch imaging component, formerly known as cold and hot buffer.

Comment by [20] [25/Aug/17]
[19], thanks for starting this page. I think you are right, from what I've seen in the new views, the requirements on the buffer are getting quite convoluted with all the services running against it. I'm a bit concerned that we seem to trying to stick to the original baseline design of the processing platform, but significantly changing the architecture above it, without taking those changes into account at all. With the work we are doing with Nyriad I think we could design a quite different approach for the buffer, which would meet the requirements at least for the bulk data. As long as we don't have concrete requirement numbers for the other services it is pretty hard to say something there. Some of these services might produce a pretty high IOPS load, which might require a dramatically different hardware setup. We will order a quite nice machine soon to test all this and that machine should actually allow us to capture about 12-15GB/s and a total capacity of around 500 TB. With that machine we should be able to prototype the design at full scale for about 1000s, provided that we can produce the data at that speed.
Comment by [2] [31/Aug/17]
[19] to update title, risk cause and event, for next sprint.
Comment by [2] [23/Oct/17]
Rescored RRB 23 Oct 17 - rescored to Major impact and Likely likelihood, because of USA Exascale performance requirements in the early 2020s. Therefore we're likely to be able to procure a suitable system at our price point.
Comment by [4] (Inactive) [03/Jan/18]
Being close to the next Risk Review Board meeting, I wonder if you have anything to state or to share that could facilitate the progress on the risk management or could better describe the status quo to the watchers. Thank you.
Comment by [19] [06/Jan/18]
On the design side [22] has provided feedback to [29] on the Buffer Model iPython notebook ~~TSK-1604~~. In Sprint 2018A this should be extended to include anticipated speeds and feeds for likely hardware available in timeframe. On the prototyping side P3 disaggregated (SAS and NvME) storage is configurable via ansible-playbooks for BeeGFS and will be completed in 2018A to support GlusterFS. The ingest network prototyping is still waiting for final cable-adaptors although the latest firmware from Mellanox appears to now support 25GbE. ~~TSK-1481~~
Comment by [2] [08/Jan/18]
We're doing a lot of work; in this sprint we need to get outcomes, and stories that are designed to produce outcomes. (Memos - stating whether or not there is a problem.) Our level of concern is increasing. However, it may resolve quickly in this sprint, with a focussed task. We will watch this, and review at the next RRB.
Comment by [2] [21/Mar/18]
Peter Braam has produced a memo. This needs review. There has been other work, but until it's reviewed, we can't rescore this risk.
Comment by [19] [21/Mar/18]
I am speaking with Peter later today and there is also the memo capturing the buffer EPIC. Both of these will contribute to this.
Comment by [2] [26/Apr/18]
Prototyping has been started to see whether we have severe issues with the buffer. This is a high priority. Until we know the outcome of the prototyping, we cannot evaluate our success in treating the

risk. Either we will have a major reworking of our architecture, or this risk will be retired.
Comment by [7] [22/Jun/18]
At next update consider input from https://jira.skatelescope.org/browse/SDPPRECDR-66.
Comment by [2] [29/Jun/18]
If we can't get the performance, the architecture may need to be reworked, or the contingency may need to be used.
Comment by [6] [12/Sep/18]
The memo on buffer prototyping work should provide information to update this risk. [21] could provide input based on the buffer prototyping memo.
Comment by [7] [19/Oct/18]
The SKA1 SDP Platform Performance Prototype (P3-AlaSKA) Prototyping report has been completed for CDR and indicates a measured performance of a prototyped system matching the SDP which is close to the required performance and the expectation is that the normal improvements provided by suppliers will meet SDP needs before 2025.
Comment by [7] [19/Oct/18]
We will continue to benchmark against available hardware. See hardware mitigation plan in cost basis of estimate.

[SDPRISK-369] Parametric Model does not account for data redistribution on processing platform Created: 19/Jul/16 Updated: 19/Oct/18	
Status:	Treated
Project:	SDP Risk Register
Component/s:	C1 Data Processor

Type:	Risk		
Reporter:	[2]	**Assignee:**	[6]
Labels:	risks		

Issue Links:	**Relates**		
	relates to	TSK-370	Performance model In Progress
Impact:	4 Major		
Probability:	3 Likely		
Exposure:	High		
Treatment:	Mitigate		
Treatment plan:	Update parametric model (and cost model) to include the effect of data redistribution.		
Residual Impact:	4 Major		
Residual Probability:	2 Low Likelihood		
Residual Exposure:	Medium		
Mitigation Start Date:	06/Jan/16		
Risk Event:	Major reduction in SDP system performance for specific workflows leading to slower processing and changes in the science observation schedule. The specific impact on the observation schedule will depend on the specific workflow that was underestimated by the Parametric Model. No impact to cost or schedule.		
Risk Cause:	Parametric model and cost model does not take Data Re-distribution over the processor platform into account, therefore cost and performance is underestimated. This is an expensive operation. This includes any re-ordering of data distribution, i.e. when changing from time-slices to frequency slices. For optimal efficiency of Pipeline runs it may be necessary to redistribute the data over the processing platform.		

Comments

Comment by [2] [28/Sep/16]
Mitigated by planning, as this is covered by the co-design activities. Suggest time delay reduce to 3, and likelihood to 3. Score 12.
Comment by [15] (Inactive) [24/Nov/16]
"Identified in the d-PDR review. "
Comment by [2] [24/Nov/16]
Risk Cause: Parametric model and cost model does not take Data Re-distribution over the processor platform into account, therefore cost and performance is underestimated. This is an expensive operation. This includes any re-ordering of data distribution, i.e. when changing from time-slices to frequency slices. For optimal efficiency of Pipeline runs it may be necessary to redistribute the data over the processing platform.Risk Event:Major reduction in SDP system performance leading to reduced science output as processing is slower.Comment: is this the 'corner-turn' issue, when changing from time-slices to frequency slices? yes - this and other data redistributions.NO MITIGATION PROPOSEDScore: 12 (high) 19/07/2016(may wish to assign onwards)
Comment by [2] [24/Nov/16]
Time score: 3.0
Comment by [2] [24/Nov/16]
Fin score: 4.0
Comment by [2] [24/Nov/16]
Safety score: 1.0
Comment by [2] [24/Nov/16]
Risk Number: 103.0
Comment by [2] [24/Nov/16]
Reputation score: 1.0
Comment by [2] [24/Nov/16]
performance score: 4.0
Comment by [2] [28/Feb/17]
Significant work has been done on the parametric model to allow consideration of data distribution. More work is under way. The parametric model can now produce workflows for possible execution frameworks based on realistic data distribution requirements and data access requirements. title: parametric model does not account for data redistribution on processing platform likelihood: low likelihood.
Comment by [2] [08/Jan/18]
Link to the work on parametric model development for 2018A. (this risk may not be explicitly looked at, but we may be able to rescore this risk in the light of that work.)
Comment by [6] [12/Sep/18]
[22], is the likelihood and impact of this risk still accurate. Could you provide updated information on this risk.
Comment by [22] [12/Sep/18]
Looks accurate. Might want to link up to TSK-406 (mostly done) and the follow-up on the (in my view)

corresponding OAR, TSK-2449.
Comment by [7] [19/Oct/18]
This is an activity the SKAO may wish to place into bridging phase 2 (i.e. to address after system CDR).

[SDPRISK-400] Hardware specifications do not reflect actual needs Created: 16/Sep/16 Updated: 19/Oct/18	
Status:	Treated
Project:	SDP Risk Register
Component/s:	C Science Data Processor

Type:	Risk		
Reporter:	[2]	**Assignee:**	[19]
Labels:	Construction, Processor-platform, risks		

Impact:	4 Major
Probability:	3 Likely
Exposure:	High
Treatment:	Mitigate
Treatment plan:	Detailed system performance analysis & modelling and benchmarking during construction.
Residual Impact:	4 Major
Residual Probability:	2 Low Likelihood
Residual Exposure:	Medium
Risk Event:	Certain algorithms require more resources (e.g. memory bandwidth) than currently estimated leading to slower processing and therefore changes in the science observation schedule. The specific impact on the observation schedule will depend on the specific algorithm that was underestimated. No impact to cost or schedule.
Risk Cause:	Hardware procured does not reflect actual needs (e.g. memory bandwidth) of algorithms (specifications derived from algorithm performance is incomplete or not accurate or wrong). This could be caused by Design Equations (which relate required algorithm performance to hardware specification) being wrong or producing the wrong specification.

Comments

Comment by [2] [24/Nov/16]
Risk Cause:Hardware specifications/requirements do not reflect actual needs (e.g. memory bandwidth) of algorithms (specifications derived from algorithm performance is incomplete or not accurate or wrong). This leads to Design Equations (which relate required algorithm performance to hardware specification) being wrong or producing the wrong specification.Risk Event:SDP does not deliver

required performance (L1 requirements). Design rework, cost and schedule overrun.Risk Register:https://docs.google.com/a/ska-sdp.org/spreadsheets/d/1cpV8yHRUh0gRnCXBsGMlF-fTMZw632qB9uhnpoXtDHI/edit?usp=sharing
Comment by [2] [24/Nov/16]
Time score: 3.0
Comment by [2] [24/Nov/16]
Fin score: 4.0
Comment by [2] [24/Nov/16]
Safety score: 1.0
Comment by [2] [24/Nov/16]
76.0
Comment by [2] [24/Nov/16]
Reputation score: 2.0
Comment by [2] [24/Nov/16]
performance score: 4.0
Comment by [7] [28/Feb/17]
Rename: Hardware specifications do not reflect actual needs [Requirements implies something else]. SDP does not deliver required performance so there are no L1 requirements on performance. We deduce them from our modelling. Mitigation: Detailed system performance modelling. Comment: Performance model developed to address this specific need. Q3/Q4 2016 performance model was updated based on PW memo. High degree of confidence that we are memory bandwidth limited and therefore reducing this risk. Impact: Major. (performance) Probability: 2 Residual risk: We do not have a hardware model. This could be mitigated by co-design.
Comment by [2] [23/Oct/17]
ASsignee to review risk event and risk cause
Comment by [2] [08/Jan/18]
there is clearly a risk about whether we're speccing the hardware correctly (e.g. enough memory on node etc...) however, it does need a little light review. What do we mean by co-design? Are we advocating co-design during construction? Or would the SKA project collaboration with cERN and argonne mitigate this? Or does our P3 work mitigate this? We need to reinvigorate this work, and make sure the design equations for hardware are properly captured in our architecture.
Comment by [2] [21/Mar/18]
No change.
Comment by [2] [29/Jun/18]
Would like to work on this in 2018D.
Comment by [7] [19/Oct/18]

Even with further vertical prototyping in 2018D and E, this risk remains with the risk scoring 2 and 4.

[SDPRISK-401] **Actual system performance worse than modelled** Created: 16/Sep/16 Updated: 19/Oct/18	
Status:	Treated
Project:	SDP Risk Register
Component/s:	C.1.2 Processor Software

Type:	Risk		
Reporter:	[6]	**Assignee:**	[5]
Labels:	Construction, risks		

Impact:	4 Major
Probability:	3 Likely
Exposure:	High
Treatment:	Mitigate
Residual Impact:	3 Moderate
Residual Probability:	3 Likely
Residual Exposure:	Medium
Risk Event:	Certain workflows perform worse than currently estimated leading to changes in the science observation schedule. The specific impact on the observation schedule will depend on the specific workflow that was underestimated. No impact to cost or schedule. The SKA project (using SAFe) is a cost capped, variable scope project and therefore performance will be impacted rather than cost.
Risk Cause:	Actual system performance is significantly worse than modelled or prototyped performance.
Existing Mitigation:	Adoption of SAFe processes will allow prioritization of work against L1 requirements that impact science delivery. The Performance Verification Learning Milestones milestones in the SDP Construction & Verification Plan are structured to incrementally test and understand the actual performance as the system is being built in order to allow performance and scaling problems to be resolved during implementation.

Comments

Comment by [2] [24/Nov/16]
Risk Cause:Actual system performance is significantly worse than modelled or prototyped performance.Risk Event:SDP does not deliver required performance (L1 requirements). Design rework,

cost and schedule overrun.Risk Register:https://docs.google.com/a/ska-sdp.org/spreadsheets/d/1cpV8yHRUh0gRnCXBsGMlF-fTMZw632qB9uhnpoXtDHI/edit?usp=sharing
Comment by [2] [24/Nov/16]
Time score: 3.0
Comment by [2] [24/Nov/16]
Fin score: 4.0
Comment by [2] [24/Nov/16]
Safety score: 1.0
Comment by [2] [24/Nov/16]
77.0
Comment by [2] [24/Nov/16]
Reputation score: 2.0
Comment by [2] [24/Nov/16]
performance score: 4.0
Comment by [7] [28/Feb/17]
This has to be an ongoing activity through until when the architecture of the hardware to be procured is clear and benchmarking is possible. The scoring is correct.
Comment by [2] [29/Jun/18]
Added to SKA Risk Register as SKA-261
Comment by [18] [02/Jul/18]
[2] and [7], please notice that the SKA-261 risks refers tot he loss of expertise between Element CDRs and beginning of SKA Construction.
Comment by [7] [19/Oct/18]
We cannot test at scale until we have a system that allows such testing. In the construction plan this is addressed explicitly.
Comment by [18] [19/Oct/18]
I understand this is exactly the core of the risk. The way you propose to handle (and mitigate) this aspect is the valuable information, together with the expected improvement of the situation once the mitigation action has been implemented.

[SDPRISK-408] HPSO experiments may not be representative Created: 16/Sep/16 Updated: 19/Oct/18	
Status:	Analyzed
Project:	SDP Risk Register
Component/s:	C1 Data Processor, C2 Delivery System, C4 Preservation System

Type:	Risk		
Reporter:	[27]	**Assignee:**	[1]
Labels:	Construction, PIP, risks		

Impact:	3 Moderate
Probability:	2 Low Likelihood
Exposure:	Medium
Treatment plan:	Mitigation: SKAO RT to analyse these related issues and provide updated inputs and changes to the Parametric Model & System Sizing.
Residual Exposure:	Medium
Risk Event:	Certain observations require more resources (compute) than currently estimated leading to changes in the science observation schedule. The specific impact on the observation schedule will depend on the specific observation that was underestimated by the Parametric Model. No impact to cost or schedule.
Risk Cause:	Operational risk: HPSOs may not be representative. We have proceeded on the basis that the HPSO experiments are representative of the kinds of experiments that SKA1 will conduct. This may not be true, in which case we will (likely) have underestimated the typical compute load for SDP.

Comments

Comment by [2] [24/Nov/16]
Risk Cause:Operational risk: HPSOs may not be representativeWe have proceeded on the basis that the HPSO experiments are representative of the kinds of experiments that SKA1 will conduct. This may not be true, in which case we will (likely) have underestimated the typical compute load for the MID instrument.Risk Event:Additional hardware required or insuffcient imaging capability.Risk Register:https://docs.google.com/a/ska-sdp.org/spreadsheets/d/1cpV8yHRUh0gRnCXBsGMlF-fTMZw632qB9uhnpoXtDHI/edit?usp=sharing
Comment by [2] [24/Nov/16]
Time score: 3.0
Comment by [2] [24/Nov/16]
Fin score: 2.0

Comment by [2] [24/Nov/16]
Safety score: 1.0
Comment by [2] [24/Nov/16]
85.0
Comment by [2] [24/Nov/16]
Reputation score: 2.0
Comment by [2] [24/Nov/16]
performance score: 3.0
Comment by [4] (Inactive) [03/Jan/18]
Being close to the next Risk Review Board meeting, I wonder if you have anything to state or to share that could facilitate the progress on the risk management or could better describe the status quo to the watchers. By the way, Happy New Year!
Comment by [6] [12/Sep/18]
This risk needs a discussion with SKAO. [23] and [24] are putting together an RT to deal with this issue and other issues related to the Parametric Model and SDP System Sizing. [1] to support from SDP. Update mitigation: SKAO RT to analyse these related issues and provide updated inputs and changes to the Parametric Model & System Sizing.
Comment by [6] [12/Sep/18]
Need to update the Risk event.
Comment by [7] [19/Oct/18]
The SDP has the tools to enable this to be further analysed. That analysis needs to be taken forward by an SKAO RT which includes science input. This will need to happen in the bridging period.

[SDPRISK-311] Cost overrun due to complexity of pipeline components Created: 24/Nov/16 Updated: 19/Oct/18	
Status:	Treated
Project:	SDP Risk Register
Component/s:	C.1.2.2 Pipelines

Type:	Risk		
Reporter:	[2]	**Assignee:**	[27]
Labels:	PIP		

Impact:	3 Moderate
Probability:	4 Highly likely
Exposure:	High
Treatment:	Mitigate
Residual Impact:	2 Minor
Residual Probability:	2 Low Likelihood
Residual Exposure:	Low
Risk Event:	Use of cost contingency in this area.
Risk Cause:	Uncertainty in estimated labour resource needed and complexity of scaling the Imaging Pipeline (deconvolution, gridding and FFT).
Existing Mitigation:	Draw upon existing components. As part of the SDP Costing to date we have considered the possible reuse from pathfinders, precursors and CASA. Also (31/08/2017) using SAFe process to prioritise work during the construction phase according to scientific and operational needs. We will have working components very early on, which we will then iteratively improve. An additional domain feature team has been added to the cost estimate for the second half of construction to further reduce this risk. The current contingency (~20%) in this area covers the remaining risk impact.

Comments

Comment by [15] (Inactive) [24/Nov/16]
Cost impact (contingency) € 5,720,119Cost of line item (incl. contingency) €30,641,011.00Contingency as % of relevant category (Hardware / Software) cost 6.1%
Comment by [2] [13/Dec/16]
Fin score: 3.0
Comment by [2] [13/Dec/16]

91.0
Comment by [7] [28/Feb/17]
The mitigation was to draw upon existing component reuse. As part of the SDP costing to-date we have considered the possible reuse from pathfinders, precursors and CASA. So the probability reduced to 3. Residual risk is required to integrate existing components, especially if the authors of those components are not available and/or documentation is poor. Title should be changed to: Cost overrun due to complexity of pipeline components. Check the risk is consistent with the contingency in the cost spreadsheet (the impact may need to go up to keep our contingency the same).
Comment by [2] [08/Jan/18]
We've talked about work to look at the use of GGPUs/multicore architectures - that may have an impact on this risk, as it may allow us to rescore this risk. Examine at the end of 2018A. Link to the ticket to look at GGPU work.
Comment by [6] [20/Mar/18]
An additional domain feature team has been added to the cost estimate for the second half of construction. This may reduce the probability of this risk - discuss at RRB23.
Comment by [2] [21/Mar/18]
Reduce probability to 2 in light of new feature team.

[SDPRISK-390] Execution Framework is immature at production (low TRL) Created: 30/Jul/15 Updated: 19/Oct/18	
Status:	Treated
Project:	SDP Risk Register
Component/s:	C.1.2.3 Execution Framework

Type:	Risk		
Reporter:	[6]	**Assignee:**	[6]
Labels:	Construction, EF, Prototyping, risks		

Issue Links:				
	Relates			
	relates to	TSK-345	COTS EF: Evaluation of Spark and Spar...	Open
	relates to	TSK-340	Execution Framework - DALiuGE	In Progress
	relates to	~~TSK-8~~	T1D COTS Framework Prototype work	Closed
	relates to	~~TSK-94~~	Generalised Execution Framework Archi...	Resolved
	relates to	TSK-1595	DALiuGE prototype testing	In Progress
	relates to	~~TSK-1753~~	Execution Frameworks upward interfaci...	Closed

Impact:	5 Catastrophic
Probability:	5 Near Certainty
Exposure:	Extreme
Treatment:	Mitigate
Treatment plan:	Architect to support multiple execution frameworks. Prototyping and analysis to identify existing frameworks by CDR which can be used at construction even if they are not optimal for the SKA requirements with a subsequent performance impact.
Residual Impact:	4 Major
Residual	1 Not Likely

Probability:	
Residual Exposure:	Low
Risk Event:	It is not possible to develop the component to provide the required performance within the budgeted cost and schedule. This is covered by the cost contingency.
Risk Cause:	The available Execution Frameworks are either not sufficiently mature at the start of construction or our understanding of the framework is immature, so that effort is wasted trying to improve its performance during construction.

Comments

Comment by [9] (Inactive) [13/Apr/16]
Hello. This ticket has passed its planned end date and is still open. Please can you update the ticket (change planned end date to in the future of change status) or let me know why it hasn't been resolved. Thanks.
Comment by [2] [20/Jul/16]
tiem delay upped from 3 to 4likelihood still 4. We have mitigation- prototyping the design. We are proposing a set of tests and stress tests to evaluateScore: 16
Comment by [5] [21/Jul/16]
Just to check : assigned to me for further clarification? Or for approving the re-score?
Comment by [5] [22/Jul/16]
I think we should replace the risk cause in this risk by identifying the actual technical risks. If I understood correctly, a low TRL is essentially a management risk as it highlights a products that are unproven and therefore require indepth risk analysis and mitigation. Having this risk in the technical section sounds to me like "we have a risk that we will have risks". The fundamental technical risks topics I see with the execution engine are:* Unclear division of functionality between the execution engine and other parts of the system means that it is not possible to integrate the execution engine with the rest of the system and therefore no working SDP can be achieved* Incomplete or incorrect requirements (especially performance) mean that an execution engine meets its performance requirements but the overall system still fails to perform adequately* The execution engine design can not meet its requirements* Complexity of the execution engine design means that it can not be implemented within the identified time frameThe latter two are execution engine risks while the top two are architectural level.We need to work through these in quite a lot more detail .
Comment by [5] [22/Jul/16]
I've made a suggestion about this risk although may be too late now to include this in the next milestone
Comment by [2] [26/Jul/16]
This is assigned to you as Project Engineer. For highly risky key items in the product tree, you 'own' the risk. One of the next steps on the RRB is to set out next steps in managing the risk (this is likely to include items like 'fill in an A3 risks template', plan future work, assign future work - these are likely to happen after the WBS meeting next week).
Comment by [20] [14/Sep/16]
The Description of this ticket is wrongly referring to SDPR-M7-G03 (0062). I would suggest to copy both the event and the cause into the description of the risk tickets or (even better) track them completely in Jira.

Comment by [2] [24/Nov/16]
Cause: Event:The fundamental technical risks topics I see with the execution engine are:1) Unclear division of functionality between the execution engine and other parts of the system means that it is not possible to integrate the execution engine with the rest of the system and therefore no working SDP can be achieved.2) Incomplete or incorrect requirements (especially performance) mean that an execution engine meets its performance requirements but the overall system still fails to perform adequately.3) The execution engine design can not meet its requirements.4) Complexity of the execution engine design means that it can not be implemented within the identified time frameThe latter two are execution engine risks while the top two are architectural level.SDP Risks
Comment by [2] [24/Nov/16]
Time score: 4.0
Comment by [2] [24/Nov/16]
Fin score: 4.0
Comment by [2] [24/Nov/16]
Safety score: 1.0
Comment by [2] [24/Nov/16]
65.0
Comment by [2] [24/Nov/16]
Reputation score: 3.0
Comment by [2] [24/Nov/16]
performance score: 3.0
Comment by [2] [28/Feb/17]
Execution framework has high uncertainty on what it will be (will it be internally designed and written, currently our baseline, with low TRL; or COTS solution where we need to identify and assess the work needed to add functionality for SDP). Mitigation: work is happening to do this. No product identified which can do this. Our prototyping is still early stage. We could still enter construction with this being the case. Develop DaLiuGE to add required functionality to increase TRL, investigate COTS options. likelihood: 80% - so near certainty. impact: catastrophic (costs) Alter title: EF is immature at production. (low TRL) If we have a low TRL solution, we are carrying risk in this area into construction. This is reflected in our costings (100% contingency). This is not a universally held view; it may be that the MeerKAT framework is more mature; this will be assessed.
Comment by [20] [01/Mar/17]
I guess it would be good to start an assessment of what functionality is still missing from DALiuGE and what that would cost to implement. Then we could start a similar assessment about one or a few other frameworks about the costs involved to get real pipelines and pipeline modules migrated and running. Note that we had been working on this since a couple of years now, thus I would not call this early stages anymore. DALiuGE is already using a whole set of existing pipeline modules from various sources and integrating them reportedly (~~DATA-167~~) is quite straight forward.
Comment by [15] (Inactive) [28/Mar/17]
At best this is at TRL level 3. If not functionally complete TRL2.

A higher TRL would require the product integrated into a system

Comment by [20] [28/Mar/17]

I have a very hard time seeing the reasoning behind this assessment!
The level 4 exit criteria reads "Documented test performance demonstrating agreement with analytical predictions. Documented definition of relevant environment". That's what our reports and the DALiuGE paper and in fact our architecture are describing in our view. We ARE running DALiuGE in a number of operational environments using real data, fully integrated into data flows involving several stakeholders and systems. We HAVE shown scalability far beyond the required level by running more than 12 million Drops on Tianhe-2 and showing the scaling properties. We have also defined the relevant environment within the SDP architecture. In addition we ARE using existing pipeline components from a variety of sources and have, just in this current sprint, integrated MPI based components as well, opening the opportunity to integrate distributed SAGEcal and ASKAP pipeline components as well.
We are targeting level 5 with our next deployment on Tianhe2. This reads "End-to-end software elements implemented and interfaced with existing systems/simulations conforming to target environment. End-to-end software system, tested in relevant environment, meeting predicted performance. Operational environment performance predicted. Prototype implementations developed."

We have planned this for this current sprint, but delays in securing the required resources did not allow to finish this. The activities leading to this are described and tracked in ~~TSK-344~~. Essentially what we will run is many instances of OSKAR2 to simulate SKA1-LOW data at scale directly on the GPU nodes of Tianhe2. This will produce SPEAD packets, which will be received by already existing SPEAD-Drops. We will then run the data through a calibration and imaging pipeline and compare the results with the sky model, which had been used to drive the simulations. This will show, as far as possible, an end-to-end system, "interfaced with existing systems/simulations conforming to target environment". Since we don't have existing prototypes of SKA ready pipeline components in the SDP, we are not able to integrate any of those, but we will integrate the current state-of-the-art existing components. We also don't have a prototype LMC, which we could integrate with at this point. The control of this DALiuGE run will have to be based on SLURM (that's what's being used at Tianhe2), if an LMC prototype can interface with SLURM we are essentially ready to go. The details of a possible or prototype implementation of the interface with CSP/SADT are not yet detailed enough, but we essentially start prototyping something like this by connecting the OSKAR2 SPEAD streams to multiple SPEAD Drops. Obviously we are also not able to run DALiuGE on anything that looks anywhere close to what the "target environment" will look like. Even our current SDP hardware baseline design is just a guess, based purely on currently available hardware, to allow us to estimate costs.

I guess it is quite hard to state that all of this is just covering TRL2: "Practical application is identified but is speculative, no experimental proof or detailed analysis is available to support the conjecture. Basic properties of algorithms, representations and concepts defined. Basic principles coded. Experiments performed with synthetic data."

or TRL3: "Development of limited functionality to validate critical properties and predictions using non-integrated software components".

Actually I had been using the NASA definitions of TRLs for software, the DoD ones are even more questionable since "Unfortunately, current conceptions of software readiness were derived from hardware-readiness notions and have proven problematic in practice." This is a quote from a very interesting SEI document describing a DoD workshop triggered by the issues with applying TRLs to software (http://resources.sei.cmu.edu/asset_files/technicalreport/2010_005_001_15305.pdf). Given that I would actually suggest to completely drop the TRL assessment for SDP, which seems to have been

a fairly uninformed decision in the first place.
Comment by [2] [02/Jun/17]
RRB: we agree that TRL isn't a good measure for SDP. We now have the option of pluggable execution frameworks, which helps mitigate this risk.
Comment by [2] [02/Jun/17]
The pluggable execution framework allows us to plug in working EFs early in development. This reduces the risk. We have much better definitions of what the EFs need to do (and a much clearer understanding of how they work). EFs have an evolution path in the architecture, and there's the option of having multiple EFs in production, each with different strengths. We may need to create a new risk around performance, but this specific risk is now reduced. Impact is still Catastrophic - there is no EF that is completely ready for production. Daliuge is well advanced, but not complete. Spark is still in experimental mode. We need to get the interfaces right. 2-6 months schedule impact is still possible. Probability is now Low likelihood - we have removed the risk associated with the whole system; there is some residual risk around interfaces.
Comment by [2] [08/Jan/18]
No change.
Comment by [2] [21/Mar/18]
We are using set-based design in construction, to help reduce risk. This is no longer an architectural risk, as we have resolved the architectural issues. the risk is that it is not possible during construction to take the identified possible EFs and develop them to the performance that is required.
Comment by [6] [26/Sep/18]
Review this risk again after receiving the reports on the EF prototyping. Then decide if this risk can be closed and or new risks need to be created for anything specific from the EF prototyping reports.

[SDPRISK-395] Uncertainty in interoperation of TM and SDP scheduling Created: 16/Sep/16 Updated: 29/Oct/18	
Status:	Treated
Project:	SDP Risk Register
Component/s:	C.1.1 Processor Platform, C3 LMC

Type:	Risk		
Reporter:	[6]	**Assignee:**	[25]
Labels:	Construction, risks		

Attachments:	interaction Between TM and SDP Scheduling.png interaction Between TM and SDP Scheduling.png
Issue Links:	**Relates** relates to ~~TSK-1306~~ Scheduling- science aspects Closed relates to ~~TSK-1267~~ SIP: Simple TM emulator Closed relates to ~~TSK-1306~~ Scheduling- science aspects Closed
Impact:	4 Major
Probability:	4 Highly likely
Exposure:	High
Treatment:	Mitigate
Treatment plan:	Model a scheduler with TM and update the architecture
Residual Impact:	2 Minor
Residual Probability:	2 Low Likelihood
Residual Exposure:	Low
Risk Event:	TM unable to schedule SDP (System level risk)
Risk Cause:	TM does not understand how to schedule SDP, including commensal observing, dealing with the offset between data collection and data processing for some pipelines.

Comments

Comment by [2] [28/Sep/16]
The mitigitation is the work being done by TM and Yongxin. Can rescore after current task complete.
Comment by [2] [24/Nov/16]
Risk Cause:TM does not understand how to schedule SDP, including commensal observing, dealing with the offset between data collection and data processing for some pipelines (Malta)Risk Event:TM unable to schedule SDP (System level risk)Risk Register:https://docs.google.com/a/ska-sdp.org/spreadsheets/d/1cpV8yHRUh0gRnCXBsGMlF-fTMZw632qB9uhnpoXtDHI/edit?usp=sharing
Comment by [2] [24/Nov/16]
Time score: 4.0
Comment by [2] [24/Nov/16]
Fin score: 3.0
Comment by [2] [24/Nov/16]
Safety score: 1.0
Comment by [2] [24/Nov/16]
121.0
Comment by [2] [24/Nov/16]
Reputation score: 2.0
Comment by [2] [24/Nov/16]
performance score: 4.0
Comment by [7] [28/Feb/17]
Comments: A lot of work post-Malta done. Reflected in increased buffer size to give a loose coupling between SDP and TM. Further work require to model in more detail. The architectural description must be updated to properly reflect the change. Treatment: Model the schedule with TM and update the architecture. Rename: Uncertainty in interoperation of TM and SDP scheduling. Scoring remains until the work is further progressed. Review again after the treatment is enacted.
Comment by [25] [09/Mar/17]
Work on the scheduler has progressed. [26] has been interacting with [30] and working on the Scheduler. [26] can you please provide an update?
Comment by [26] [12/Mar/17]
[25]We are working on priority scheduling on several conditions. For example, we provide the critical path in a task DAG with high priority and perform the scheduling based on this.
Comment by [25] [15/Mar/17]
Hi [26] Could you please send me a link to documents or work aimed at the scheduler?
Comment by [26] [15/Mar/17]
[25] We are drafting updated document with updates on: 1) GUI wrapper of scheduler to enable interaction between SCHEDULER, TM as well as LMC (an initial interface can be referred at "http://wirelesslab.sjtu.edu.cn/~klh/cts/" to support variable number of parameters from TM, LMC and computation platforms); 2) enabling priority scheduling by transforming priority scheduling into preemptive scheduling.
Comment by [2] [02/Jun/17]

Rosie, Ferdl, Alan had a meeting - scheduling between TM and SDP has been agreed (though some further modelling needs to happen). It's documented in activity diagrams, and the TM ICD should be finished next week, which will describe those diagrams, and there will be specific use cases described. We still need to work out some responses to failures and so on, which means the risk is not completely avoided. There may be additional risks to come around non-compliance.
Comment by [6] [19/Oct/18]
Meeting held on 21 Sept with [23] and members of the SKA Software Architecture Team (SARCH). Resources, Capabilities and Scheduling of SDP was discussed. The understanding of SDP scheduling at the system level was discussed and clarified. The SDP-TM ICD has been updated following this meeting. Further work in this area is on the backlog of the SARCH team and this work will start in December once the SKA ART is launched.
Comment by [26] [29/Oct/18]
There have been some agreements between TM observation planner and SDP scheduler. The interaction flow was proposed in "SKA TM OBSMGT SDP PLANNING INTERFACE TECHNICAL NOTE". The flow is illustrated in the attached diagram (courtesy of Steve, Stewart and Alan).