SDP Memo FP32 vs. FP64

Document Type……...Memo
Author……..... ……..……..John Taylor
Release Date……. 2017-09-07
Document Classification…… Project Only
Status……..... Released

Lead Author	**Designation**	**Affiliation**
John Taylor		Univ of Cambridge

Table of Contents

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-TEL-SDP-0000046 SDP Costing Basis of Estimate

[AD02] SKA-TEL-SDP-0000108 Updated SDP Cost Basis of Estimate Feb 2017

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.

[RD01] https://confluence.ska-sdp.org/display/MN/NZA-SIP-P3-F2F+-+22-May-17?preview=%2F222986247%2F223805468%2FGriffin_v2.pdf

[RD02] https://images.nvidia.com/content/tesla/pdf/nvidia-tesla-p100-PCIe-datasheet.pdf

3 SCOPE

This document provides an analysis of primarily the impact on capital cost assuming that the SDP pipelines can be supported by a relaxation in the float-point accuracy from 64-bit (current baseline) to 32-bit in entirety. No comments are made on the efficacy of this relaxation in terms of the science impact. This may require further analysis. In addition to the relaxation on floating point performance scenarios are also presented in terms of increasing pool memory (DRAM per compute node) and decreasing the GPU count per node to one.

4 SUMMARY

This document provides a comparison between the Feb 2017 Cost update [AD02] for the now called delivery baseline. [AD02] was an update to the SDP Costing Basis of Estimate [AD01]

Construction WBS (CAPEX)	Feb 2017	Aug 2017		
	Cost	Scenario1	Scenario 2	Scenario 3
Hardware				
Mid	€30,725,174	€20,347,360	€20,573,425	€22,418,102
Low	€29,437,422	€18,305,778	€18,453,097	€19,900,806
Total	**€60,162,596**	**€38,653,138**	**€39,026,522**	**€42,318,908**
Cost (incl Contingency)				
Mid	*19%*			
Low	*21%*			
Scenario 1	*Replace all FP64 with FP32 for GPU only*			
Scenario 2	*As scenario 1 with double pool memory for FP32*			
Scenario 3	*As scenario 2 with 1 GPU per compute node*			

Table 1: A comparison breakdown of construction costs (CAPEX) including contingency for the hardware only between the February 2017 submission and with the Scenarios as shown above

5 MOTIVATION

Recent work performed by [RD01] has indicated that astronomic pipelines can be satisfied by using single precision floating point arithmetic as opposed to the hitherto assumption of double precision. As a consequence of that, costs are presented here for the hardware for SDP based on this assumption for the case where all pipelines are deemed to be fulfilled in single precision.

In the current system sizing all floating point capability is assumed to be contributed by the compute node accelerators: GPUs under the current hardware concept. Thus the change in system sizing can be readily satisfied by modification on the peak performance of the GPU and assuming that the extrapolated performance is identical to that achieved for the double precision case, which has been the case for previous generations of the Nvidia GPU series.

The current hardware concept is based on the latest generation Pascal P100 GPU with 5TF (double precision). For the single precision case we have assumed a peak performance of 9.7TF [RD02][1]. The 10% efficiency from previous cost submissions is maintained.

5.1 Changes from the previous costing submission (Feb 2017)

The current hardware costed concept has been modified to produce three separate scenarios from the previous costed submission of Feb 2017.

Scenario 1. The Peak GPU Performance has been modified to 9.7TF with initial conditions (start date) as per the Feb 2017 cost submission.

Scenario 2. The Pool Memory per node - $M_{cu,Pool}$ has been modified to 2x with the same conditions as Scenario 1. Thus

Pool Memory per node (Gbytes)	640	960

Scenario 3. The number of GPUs had been reduced to one with the same conditions as Scenario 2.

The following Table provides the base cost (excluding contingency) for the SKA-LOW case.

Costs	Feb 2017	Aug 2017		
	Cost	Scenario1	Scenario 2	Scenario 3
Low				
Compute	€18,555,044	€10,227,018	€10,343,293	€10,995,535
Buffer	€2,054,112	€2,051,586	€2,051,586	€2,054,112
Pool Memory	€282,012	€145,345	€290,689	€581,378
LLN	€503,991	€281,947	€281,947	€518,089
BDN	€420,886	€281,559	€281,559	€428,443
PDN	€109,905	€61,048	€61,048	€112,978
Management	€21,299	€21,299	€21,299	€21,299
Infrastructure & Racks	€162,045	€94,380	€94,380	€167,250
Total	**€22,109,295**	**€13,164,181**	**€13,425,801**	**€14,879,084**

Table 2 A comparison of the costs (excluding contingency) for SDP-Low under 3 different scenarios as explained above

[1] [RD02] Also indicates that the double precision performance is 4.7TF and the current hardware costed concept should be updated to this figure.

In addition to the costs, the Table below provides some useful parameters on the effect of ingest on the Buffer function per Compute Node and per Compute Island per Scenario. In the above analysis the effect of the receive function has not been assessed in terms of the number of 2nd stage network switches.

Buffer Parameters	Feb 2017	Aug 2017			
		Scenario1	Scenario 2	Scenario 3	Units
Low					
Total Ingest rate	0.39				TB/s
Buffer Size	50.60				PB
Buffer storage per node	32.44	62.94	62.94	32.44	TB
Buffer storage per island	1.74	3.16	3.16	1.74	PB
Ingest rate per node	1.99	3.87	3.87	1.99	Gb/s
Under-subscription at node	12.54	6.46	6.46	12.54	
Ingest rate per Compute Island	107.24	194.38	194.38	107.24	Gb/s

Table 3 A comparison of the key parameters for the receive function for SDP-Low under 3 different scenarios as explained above