The Run II DST

The Data Tier Committee[1]

October 13, 2000

The run II DST will be the default RECO output and will be stored in the tape robot[2]. A subset will be available on disk at all times, possibly allowing consistent access to the full sample via disk using a freight train[3]. We intend the DST to allow complete analysis at the physics object level, with some additional intermediate information from event reconstruction, and that full event reconstruction be possible from the DST. Roughly speaking, the DST will be a slight superset of the TMB plus packed raw data.

Table 1 shows the predicted sizes of each component on the DST. The contents and event sizes are from estimates presented to the data tier committee[4], based primarily on $t\bar{t}+2$ interaction events processed through standard DØGSTAR, digitization and reconstruction. For this sample, the predicted size is 105 kB/evt. However, electronic and beam noise is not part of the Monte Carlo generation, so we have also estimated the size when the SMT occupancy is arbitrarily increased to 5%. For this case, the DST becomes 150 kB/evt.

Data Type	Size (kB/evt)
Physics Objects	7
Vertices	8
Minimal Gtrack	4
Cal. clusters	9
MUON	4
SMT raw data	20(65)
CFT " "	18
CPS " "	5
FPS " "	1
CAL " "	22
TMB chunk	7
TOTAL	105(150)

Table 1: The size of each component of the DST base on $t\bar{t}+2$ interaction events. The numbers in parentheses are for the case in which the SMT is arbitrarily forced to have 5% occupancy.

The physics objects, vertices and calorimeter clusters are expected to be the standard RECO chunks without format changes. The *gtrack* information on the DST is a subset of the standard *GTrack* data. We retain the fit results at the first and last layer for all fits but not the intermediate layer information. Furthermore, the individual hits will not be kept, but must be redetermined if needed. This form of the *gtrack* fills the requirement of the offline vertex–finding algorithms as well as providing slightly

more detailed tracking information than is available on the TMB. The calorimeter clusters are needed to link jets and the constituent calorimeter cells, and because the muon data is a minor component it includes the hits, segments and tracks.

The raw data on the DST may be stored in a different format than the raw data chunk itself, probably in separate chunks for each detector. In all cases we imagine uncalibrated data is stored and no loss of precision occurs. The calorimeter information will be packed as in the raw data chunk. The SVX–based(SMT, CFT, CPS, FPS and FPD) raw data can be packed using a loss–free method in which overall SVX chip–addresses are packed as bit masks, signficantly reducing the required space(See below).

Not included in the size estimate are general event information, for example run and event number and instantaneous luminosity, and the trigger results. Both of these categories must be present on the DST. The size estimate does not include any reduction occuring from the standard bit–stream compression used when writing in dØ om EVPACK format.

The standard TMB is designed to be unpacked back into the physics objects transparently to TMB users. Because the TMB stores variables with reduced precision, this could lead to trivial numerical differences caused by roundoff (*e.g.* bin migration) between analyses performed from the DST and from the TMB. The committee recommends that the DST physics objects are retained, but that by default, they are superceded by the equivalent unpacked TMB objects when event data is read. This insures that unless special action is taken, analyses performed on DST and TMB will have identical numerical results. This substitution should be selected via RCP in order to permit use of the original physics objects if desired.

If the clusters and hits used in *e.g.* the track fitting are to be reconstituted to reproduce results identical to the original reconstruction, all chunk history must be kept, and statically–linked back versions of RECO must be routinely available. For general rereconstruction, this information is not needed.

SVX data packing scheme

The raw data for detectors using SVX read out is packed in a format such that address size can be minimized. The data for a given is unpacked by scanning the full packed record for that subdetector. This scheme also implies (arbitrary) static look–up–tables converting hardware addresses to SVX indices and vice-versa. For all detectors this reduces the size significantly.[5] The packing scheme is illustrated below.

The algorithm requires as the starting point a bit mask with one bit for each SVX in a given subdetector. The SVX chips in a given subdetector are numbered from $1 - N_{SVX}$, and if the ith bit in the bit mask is set, then the corresponding SVX chip has at least one channel with data. If the bit is clear, the chip has no hits. The channel data are then stored in one of two formats. The format should be chosen on an event-by-event basis for each subdetector. The format resulting in the smallest size is chosen, and a single bit is used to indicate the format.

In the first format, SVX data for each chip with hit channels are represented by a 128 bit bitmask, with a set bit implying the corresponding channel was read out. The

bitmask is followed by the unmodified 8–bit SVX pulse heights in channel order. Pulse heights are present only for channels with a signal. The data for successive SVX–chips follow one another in the (bitmask, pulse heights) format with no additional indexing.

The previous format results in the smallest size as long as a given SVX chip has at least 15 hit channels. For the case in which less than 15 channels are read out, the SVX–chip bitmask is not used, and instead, the data for a single channel is represented using an 8–bit channel offset and the 8–bit pulse height. Successive channels are packed back-to-back with a count of the number of hit strips starting the record for each SVX chip.

Packing Method #1 (exampe having 2 SVX chips w/hits)

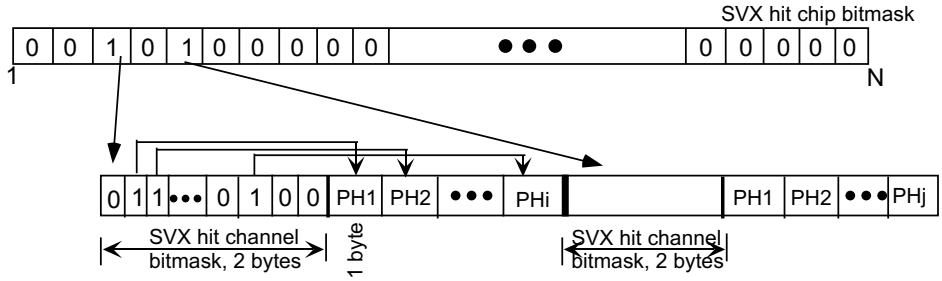

Packing Method #2

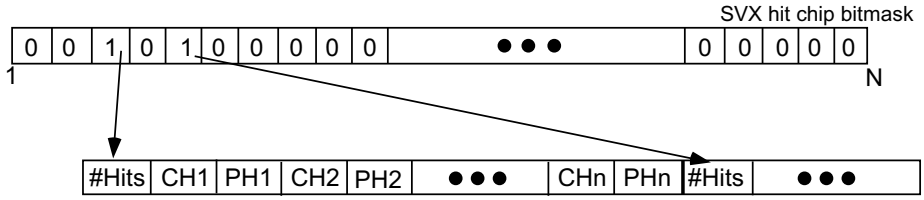

PH = 8 bit pulse height
CH = 8 bit channel offset

[1] U. Heintz, J. Hobbs, B. Klima, J. Krane, L. Lueking, S. Protopopescu, M. Strovink, A. Zylberstejn
[2]The DST size prohibits storing the complete DST on disk.
[3] See DØ Note 3313, *New Phenomena Comments on the DØ Run II Data Access.*
[4] Presentations regarding DST contents and event size were given by Dave Adams, Ela Barberis, Don Coppage, Mike Hildreth, Slava Kulik, Meena Narain, Serban Protoposcu, Dennis Shpakov and Daria Zieminska.
[5] For CFT, this improvement is slightly exceeded by a packing scheme which reduces the precision of the digitized data with negligible loss in tracking performance. If we need to eke out every possible byte, this scheme could be used instead. *Mike Hildreth to the data tier committee*