

NGON modification proposal

Unstructured grid connectivity

Author: Pierre-Jacques Legay, ONERA

Contact: pierre-jacques.legay@onera.fr

Index table

1 Rationale for the modification proposal.....	1
1.1 Current SIDS description for unstructured grid connectivity.....	1
1.2 Limitations for NGON_N and MIXED.....	3
2 NGON modification proposal.....	4
2.1 Solution description.....	5
2.2 Solution analysis.....	6
3 Example of extension.....	9
4 Implementation note.....	10
5 Conclusions.....	10
6 Appendix: Document modification list.....	10

This CPEX focuses on the NGON representation. The rationale for requiring an extension to CGNS/SIDS for unstructured grid connectivity is detailed in the first part of this document. The second part details the proposal which includes a solution to the problem and an impact analysis of that solution on the SIDS and on the performance.

1 Rationale for the modification proposal

The first section is a reminder of the current SIDS description for unstructured grid connectivity. The following section details the problems induced by this description.

1.1 Current SIDS description for unstructured grid connectivity

The unstructured grid connectivity is stored in the `Elements_t` of the CGNS/SIDS. As described hereafter, its storage depends on the elements type:

For all element types except MIXED, NGON_n, and NFACE_n, `ElementConnectivity` contains the list of nodes for each element. If the elements are sorted, then it must first list the connectivity of the boundary elements, then that of the interior elements.

```
ElementConnectivity = Node11, Node21, ... NodeN1,  
                    Node12, Node22, ... NodeN2,  
                    ...  
                    Node1M, Node2M, ... NodeNM
```

where M is the total number of elements (i.e., `ElementSize`), and N is the number of nodes per element.

`ElementDataSize` indicates the total size (number of integers) of the array `ElementConnectivity`. For all element types except MIXED, NGON_n, and NFACE_n, the `ElementDataSize` is given by:

```
ElementDataSize = ElementSize * NPE[ElementType]
```

where `NPE[ElementType]` is a function returning the number of nodes for the given `ElementType`. For example, `NPE[HEXA_8]=8`.

32 When the section `ElementType` is `MIXED`, the data array `ElementConnectivity` contains
 33 one extra integer per element, to hold each individual element type:
 34 `ElementConnectivity = Etype1, Node11, Node21, ... NodeN1,`
 35 `Etype2, Node12, Node22, ... NodeN2,`
 36 `...`
 37 `EtypeM, Node1M, Node2M, ... NodeNM`
 38 where again `M` is the total number of elements, and `Ni` is the number of nodes in
 39 element `i`. In the case of `MIXED` element section, `ElementDataSize` is given by:
 40 `ElementDataSize =sum(n=[start, end], NPE[ElementTypen] + 1)`

41 Arbitrary polyhedral elements may be defined using the `NGONn` and `NFACEn`
 42 element types. The `NGONn` element type is used to specify all the faces in the grid,
 43 and the `NFACEn` element type is then used to define the polyhedral elements as a
 44 collection of these faces.

45 I.e., for `NGONn`, the data array `ElementConnectivity` contains a list of nodes making
 46 up each face in the grid, with the first value for each face defining the number of
 47 nodes making up that face:
 48 `ElementConnectivity = Nnodes1, Node11, Node21, ... NodeN1,`
 49 `Nnodes2, Node12, Node22, ... NodeN2,`
 50 `...`
 51 `NnodesM, Node1M, Node2M, ... NodeNM`
 52 where here `M` is the total number of faces, and `Ni` is the number of nodes in face `i`. The
 53 `ElementDataSize` is the total number of nodes defining all the faces, plus one value
 54 per face specifying the number of nodes making up that face.

55 Then for `NFACEn`, `ElementConnectivity` contains the list of face elements making up
 56 each polyhedral element, with the first value for each polyhedra defining the number
 57 of faces making up that polyhedral element.
 58 `ElementConnectivity = Nfaces1, Face11, Face21, ... FaceN1,`
 59 `Nfaces2, Face12, Face22, ... FaceN2,`
 60 `...`
 61 `NfacesM, Face1M, Face2M, ... FaceNm`
 62 where now `M` is the total number of polyhedral elements, and `Ni` is the number of faces
 63 in element `i`. The sign of the face number determines its orientation. If the face
 64 number is positive, the face normal is directed outward; if it's negative, the face
 65 normal is directed inward.
 66 `ElementDataSize =sum(n=[start, end], FPP[ElementTypen] + 1)`
 67 where `FPP[ElementTypen]` is a function returning the number of faces per polyhedra.

68 The following figure is set up to ease comprehension and comparison in the following sections:

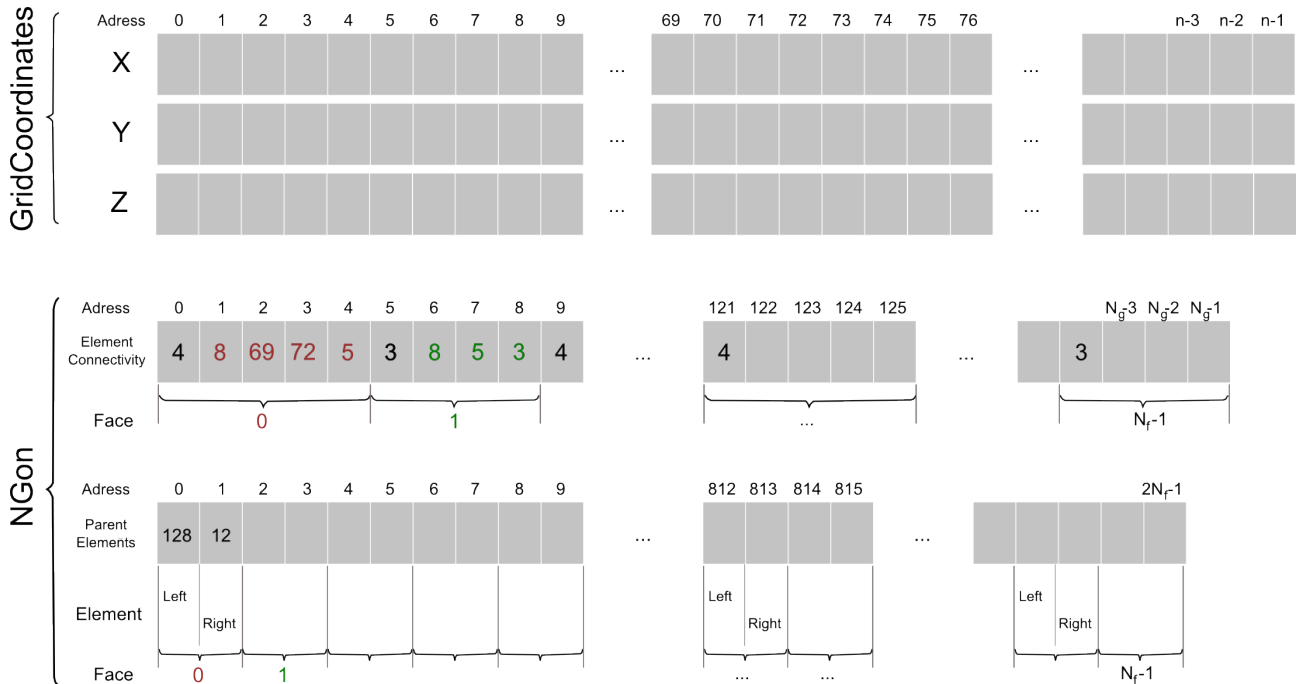



Illustration 1: CGNS/NGON current face based representation.

69 This figure shows an unstructured element using the NGON representation. It describes the physical
 70 construction of the CGNS related arrays GridCoordinates, ElementConnectivity, and
 71 ParentElements.

72 Notation:

- 73 • N_g : represents a vertex from the current face.
- 74 • N_f : represents a face of the current element.

75 1.2 Limitations for NGON_N and MIXED

76 The ElementConnectivity array can be very large for industrial CFD case. As a direct consequence
 77 we should be able to load this array fully or partially on multiple threads.

78 In the above description of the CGNS/NGON face based representation, the ElementConnectivity
 79 array mixes two data types which are interdependent :

- 80 • the number of nodes for each face : N_{nodes1}
- 81 • a list of nodes for a face : $Node_{11}, Node_{21}, \dots, Node_{N1}$

82 This implies that we cannot efficiently split this array to be read in parallel.




Illustration 2: ElementConnectivity interlaced data.

83 In parallel, the interlaced data requires that the entire ElementConnectivity array be read on every
 84 processor. This has a significant impact on IO performance.

85 NB: As described in the SIDS, the representation of MIXED and NFACE elements is identical to
 86 the NGON representation. As such, the parallel read suffers from the same performance issues. The
 87 issues of the MIXED and NFACE elements shall be addressed in a later CPEX.

88 2 NGON modification proposal

89 This proposal modifies the NGON SIDS representation to provide efficient parallel IO access and to
 90 optimize the data representation.

91 2.1 Solution description

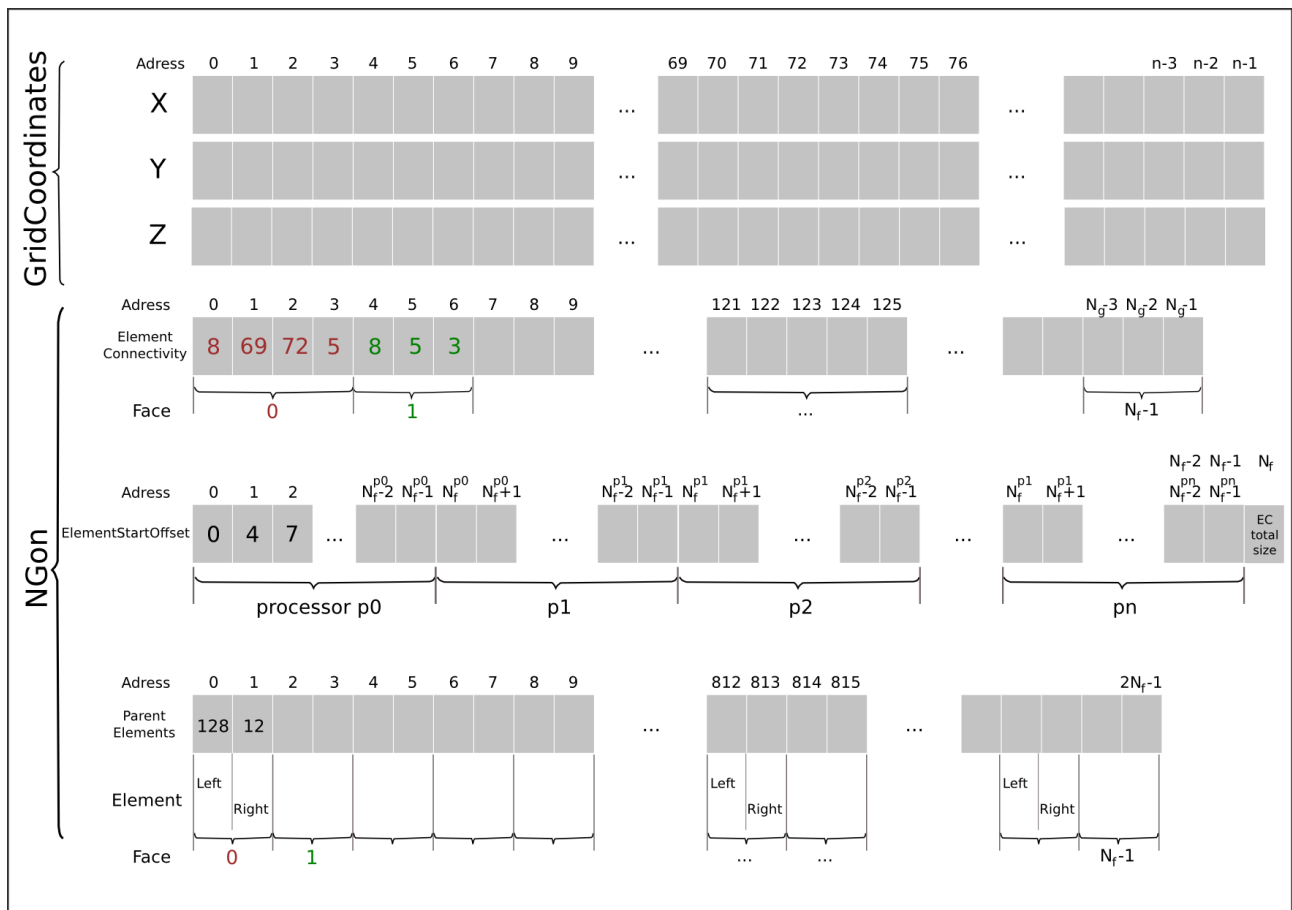


Illustration 3: NGON with new ElementConnectivity array and ElementStartOffset position array.

92 In this representation the ElementConnectivity array is deinterlaced. Illustration 4 compares the
 93 current standard versus the new deinterlaced ElementConnectivity array:




Illustration 4: Current ElementConnectivity vs. new ElementConnectivity.

94 The face vertex count has been removed from the connectivity. As a direct consequence the new
 95 array contains a unique data type.

96 Indices needed to read or analyze the ElementConnectivity array are stored in the new
 97 ElementStartOffset array.

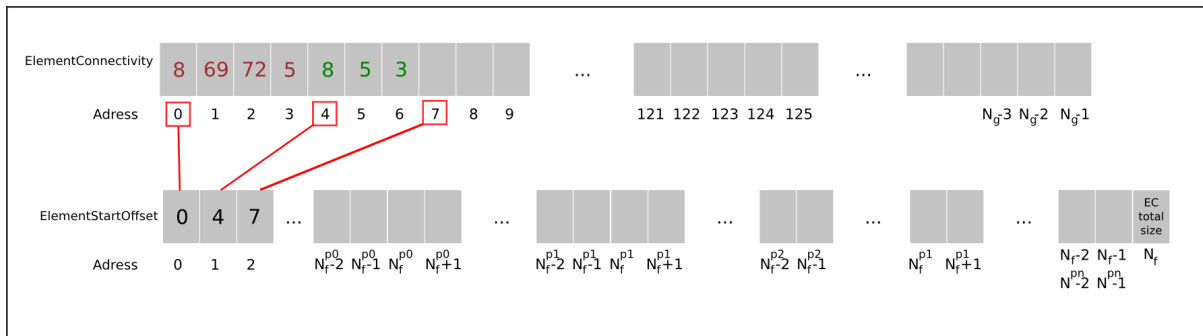


Illustration 5: ElementStartOffset lists the face position in the ElementConnectivity and its last value indicates the ElementConnectivity total size.

98 This array lists the position in the ElementConnectivity of the first vertex for each face and its last
 99 value is the ElementConnectivity array size. This CGNS node is of type DataArray_t which allows
 100 SIDS implementations to use int64 integers. For example, the Mid Level Library implementation
 101 should use the cgsizet_t type. The elementStartOffset read can be distributed between P processors.
 102 Its size is N_f+1 with N_f being the number of face in the ElementConnectivity array. The
 103 ElementStartOffset array makes it easy for a process or thread to read a portion of the
 104 ElementConnectivity array.

105 NB: The last value of the ElementStartOffset allows easy access to the last vertex of the last face of
 106 the ElementConnectivity array.

107 2.2 Solution analysis

108 2.2.1 CGNS standard modification

109 It strongly modifies the representation of the ElementConnectivity node and it inserts the
 110 new array ElementStartOffset. These modifications should be reflected in the current CFD
 111 database and CGNS related code to insure the continuity of the computational capability.
 112 See section 4 'Implementation note'.

113 2.2.2 Data consistency

114 The data type of both arrays ElementConnectivity and ElementStartOffset are consistent.

115 2.2.3 Optimal data size

116 This solution does not duplicate data thus the global data size is unchanged and stays
117 optimal.

118 2.2.4 ElementStartOffset -- an incremental index

119 In this proposal we choose to use the face position in the element connectivity instead of the
120 face number of nodes (ElementStartOffset=[0,4,7,11, ...] instead of [4, 3, 4, ...]).

121 The rationales are all IO read/write oriented and detailed here after:

122 1. Processor inter-dependency

123 To access the ElementConnectivity data we need to know the position of the face in the
124 ElementConnectivity array.

125 → “Face vertex number” solution

126 In this configuration we would need to sum over the index array to obtain the face position
127 in the ElementConnectivity array. This means either create a new table or perform the
128 computation as many times as needed.

129 Another problem is that we need the result of the sum of face vertex of processor P_0 to
130 obtain the location of the first vertex of the first face assigned to processor P_1 .

131 This behavior generalizes with the need to sum the face vertex of processors P_0, P_1, \dots, P_{M-1} to
132 obtain the location of the first vertex of the first face assigned to processor P_M .

133 **This is not a complex operation but we do not like the idea of imposing a dependency**
134 **on all previous processors computations.**

135 → ”Face position” solution

136 In this configuration the ElementStartOffset array can be split on P processors and each
137 processor related part directly gives the location of the face vertex from the
138 ElementConnectivity array.

139 2. Full vs. partial load

140 As stated above we need to know the position of the face in the ElementConnectivity array.

141 → “Face vertex number” solution

142 If we want to access data in processor P_M , then we need to access data on all processors
143 P_0, P_1, \dots, P_{M-1} . This means that we need to load the index array on every processor of inferior
144 rank.

145 → “Face position” solution

146 To access the ElementConnectivity we need the boundary of locations for processor P_M .

147 These boundaries can be partially loaded from the ElementStartOffset array by getting the
148 first element of processor P_M part and the first element of processor P_{M+1} part. These two
149 integers indicate the section of ElementConnectivity to be loaded on proc P_M (1st element of
150 P_M till 1st element of $P_{M+1}-1$).

151 NB: The last value of the ElementStartOffset allows to apply the same operator to the last
152 section of the ElementConnectivity array. For the last section, the load is performed from P_M
153 to $P_{M+1}-1$ with P_{M+1} being the ElementConnectivity size.

154 **With this proposal, we need to load only two integers from the ElementStartOffset**
155 **array on a processor to fully access the ElementConnectivity array.** Partial load for
156 CGNS/HDF5 is available in CHLone for example. CHLone (<http://chlon.sourceforge.net>)

157 is a python module implementing the SIDS-to-Python mapping which allows a simple
158 load/save of CGNS/HDF5 files.

159 3. direct access to a face connectivity

160 For some specific unstructured software, it is interesting to have a simple access to a specific
161 face in the ElementConnectivity array.

162 This is only allowed with the “Face Position” solution.

163 4. Access to face position and face number of nodes

164 Using the incremental index method, it is easy to calculate the “face number of nodes” for
165 element ‘i’ as ElementStartOffset[i+1]-ElementStartOffset[i] is an O(1) calculation. So the
166 selected method gives both data (face position and face number of nodes) in O(1) where the
167 alternate method would give face number of nodes in O(1), but face position in O(N).

168 **2.2.5 ElementStartOffset content**

169 Using an incremental index --the face position in the ElementConnectivity array-- could lead
170 to large numbers in the ElementStartOffset array. However, these numbers are limited to the
171 size of the ElementConnectivity array as they give access to addresses of that array.
172 Moreover, the cgns type used to store the ElementStartOffset node is DataArray_t which
173 allows SIDS implementations to use I8 or cgsizet types.

174 **2.2.6 ElementStartOffset last value**

175 The ElementStartOffset lists the first vertex position of each face in the
176 ElementConnectivity. Then we added the ElementConnectivity size as last value. The
177 rationale behind this is to improve data access when looping over the faces.

178 → Positions without ElementConnectivity size

179 In this configuration an iteration loop over NGONs could be written as follows:

```
180 for (i=0; i < NGON; i++) {  
181     end = (I == NGON-1) ? ElementConnectivitySize  
182     :ElementStartOffset[i+1];  
183     for (j=ElementStartOffset[i]; j < end; j++) {  
184         entry = ElementConnectivity[j];  
185         // do something here  
186     }  
187 }
```

188 For each NGON we have to check if we are reaching the last NGON to handle the access to
189 the last vertex of the last face.

190 → Positions with ElementConnectivity size (current proposal)

191 In this configuration an iteration loop over NGONs could be written as follow:

```
192 for (i=0; i < NGON; i++) {  
193     for (j=ElementStartOffset[i]; j < ElementStartOffset[i+1]; j++) {  
194         entry = ElementConnectivity[j];  
195         // do something here  
196     }  
197 }
```

198 The last value of the ElementStartOffset gives a boundary to access the last vertex of the last
199 NGON without the conditional assignment.

200 **3 Example of extension**




Illustration 6: Simple configuration demonstrating the NGON proposal.

Node Name	Dimensions	Bytes
ElementConnectivity (Old)	740	2960
ElementConnectivity (New)	1	2368

Illustration 7: Previous vs. new ElementConnectivity array

Node Name	Dimensions	Bytes
ElementStartOffset	149	1192

Illustration 8: ElementStartOffset array

201 **4 Implementation note**

202 This section aims to reflect discussions during previous CGNS meetings related to the
203 implementation impact of this modification proposal. As detailed above, this proposal strongly
204 modifies the representation of the ElementConnectivity node. As a consequence we need to address
205 the backward compatibility issue.

206 The following solutions were discussed in order to insure backward compatibility:

- 207 • CGNSLibraryVersion

208 In software, the CGNSLibraryVersion could be used to determine whether the CGNS file
209 uses the current or proposed NGON representation. This number attached to the standard
210 version will allow parsing tools to use the correct NGON representation.

- 211 • Conversion utility

212 A dedicated tool could be written to convert current CGNS files to the new NGON
213 representation

-

214 These two solutions should allow a smooth transition for all users of the CGNS standard.

215 **5 Conclusions**

216 This proposal for the NGON representation addresses the HPC issue due to the interlaced data
217 representation of the unstructured connectivity description. The solution optimizes the data
218 representation for parallel IO and has a low impact on the SIDS to deinterlace data.

219 As shown in the first part, MIXED and NFACES elements are equally concerned by the parallel IO
220 access. As such their SIDS representation should be updated accordingly to the solution chosen for
221 NGON. This issue should be addressed in a later CPEX.

222 **6 Appendix: Document modification list**

223 1. Following Marc Poinot remarks:

- 224 • Rename array from FaceConnectivityPosition to ElementStartIndex (text and figures)
- 225 • Reformulate a sentence describing the ElementStartOffset at line 101
- 226 • Remove remark concerning array size limitation in section 2.1.2.5

227 2. Following Robert Bush suggestion:

- 228 • Added section “Implementation note” to reflect CGNS committee discussions.

229 3. Remove multiple typo.

230 4. Following Gregory Sjaardema feed back:

- 231 • Modify the ElementStartOffset definition to improve the ability to loop on the data array.
232 The new size is N+1 and last index represent the ElementConnectivity total size.
233 This modification is propagated through sections 2 and 3.
- 234 • Correct multiple grammar, typographical or formatting issues.

235 5. Following Richard Hann feed back:

- 236 • Rename array from ElementStartIndex to ElementStartOffset (text and figures)
- 237 • Specify ElementStartOffset data type as cgtype_t.