

NGON modification proposals

Unstructured grid connectivity

Author: Pierre-Jacques Legay, ONERA

Contact: pierre-jacques.legay@onera.fr

Index table

1 Rationales for the modification proposals.....	1
1.1 Current SIDS description for unstructured grid connectivity.....	1
1.2 Limitations for NGON_N and MIXED.....	3
2 NGON modification proposals.....	4
2.1.1 Solution description.....	4
2.1.2 Solution analysis.....	5
3 Example of extension.....	7
4 Conclusions.....	8
5 Annexe : Document modification list.....	8

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 performances.

1 Rationales for the modification proposals

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.

33 When the section **ElementType is MIXED**, the data array ElementConnectivity
34 contains one extra integer per element, to hold each individual element type:

```
35 ElementConnectivity = Etype1, Node11, Node21, ... NodeN1,  
36                     Etype2, Node12, Node22, ... NodeN2,  
37                     ...  
38                     EtypeM, Node1M, Node2M, ... NodeNM
```

39 where again M is the total number of elements, and Ni is the number of nodes in
40 element i. In the case of MIXED element section, ElementDataSize is given by:

```
41 ElementDataSize =sum(n=[start, end], NPE[ElementTypen] + 1)
```

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

46 I.e., for **NGON_n**, the data array ElementConnectivity contains a list of nodes making
47 up each face in the grid, with the first value for each face defining the number of
48 nodes making up that face:

```
49 ElementConnectivity = Nnodes1, Node11, Node21, ... NodeN1,  
50                     Nnodes2, Node12, Node22, ... NodeN2,  
51                     ...  
52                     NnodesM, Node1M, Node2M, ... NodeNM
```

53 where here M is the total number of faces, and Ni is the number of nodes in face i. The
54 ElementDataSize is the total number of nodes defining all the faces, plus one value
55 per face specifying the number of nodes making up that face.

56 Then for **NFACE_n**, ElementConnectivity contains the list of face elements making up
57 each polyhedral element, with the first value for each polyhedra defining the number
58 of faces making up that polyhedral element.

```
59 ElementConnectivity = Nfaces1, Face11, Face21, ... FaceN1,  
60                     Nfaces2, Face12, Face22, ... FaceN2,  
61                     ...  
62                     NfacesM, Face1M, Face2M, ... FaceNM
```

63 where now M is the total number of polyhedral elements, and Ni is the number of faces
64 in element i. The sign of the face number determines its orientation. If the face
65 number is positive, the face normal is directed outward; if it's negative, the face
66 normal is directed inward.

```
67 ElementDataSize =sum(n=[start, end], FPP[ElementTypen] + 1)
```

68 where FPP[ElementTypen] is a function returning the number of faces per polyhedra.

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

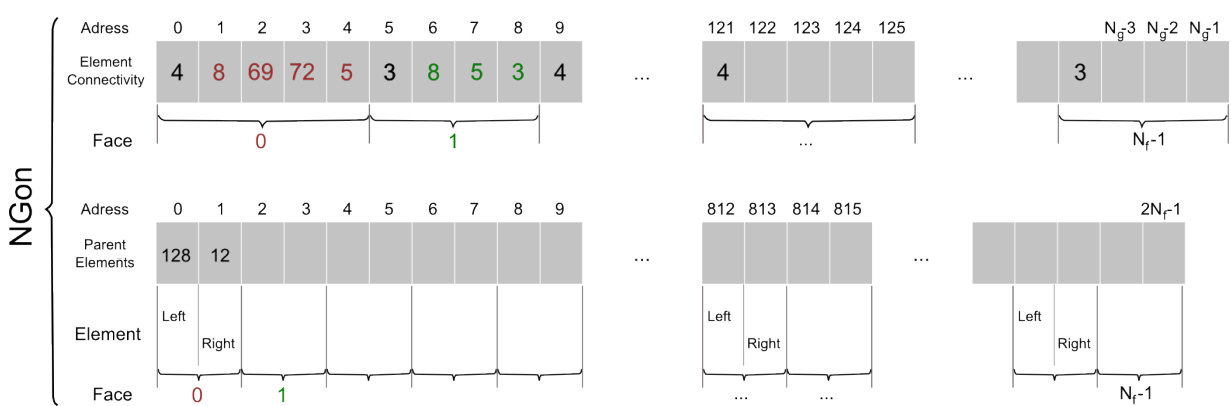




Illustration 1: CGNS/NGON current face based representation.

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

73 Notation:

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

76 1.2 Limitations for NGON_N and MIXED

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

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

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

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



Illustration 2: ElementConnectivity interlaced data.

84 In parallel, these interlaced data impose the load of the complete ElementConnectivity array on
 85 every processors. This has a significant impact on IO performances.

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

89 2 NGON modification proposals

90 This proposal offers to alter the NGON SIDS representation in order to allow an efficient parallele
 91 IO access and to optimize the data representation.

92 2.1.1 Solution description




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

93 In this representation the ElementConnectivity array is un-interlaced. The following pictures
 94 compares the current standard versus the new un-interlaced ElementConnectivity array:




Illustration 4: Current ElementConnectivity vs new ElementConnectivity.

95 The number of face vertex have been removed from the connectivity. As a direct consequence the
 96 new array includes a unique data type.

97 Adresses needed to read or analyze the ElementConnectivity array must be obtained from the new
 98 FaceConnectivityPosition array.

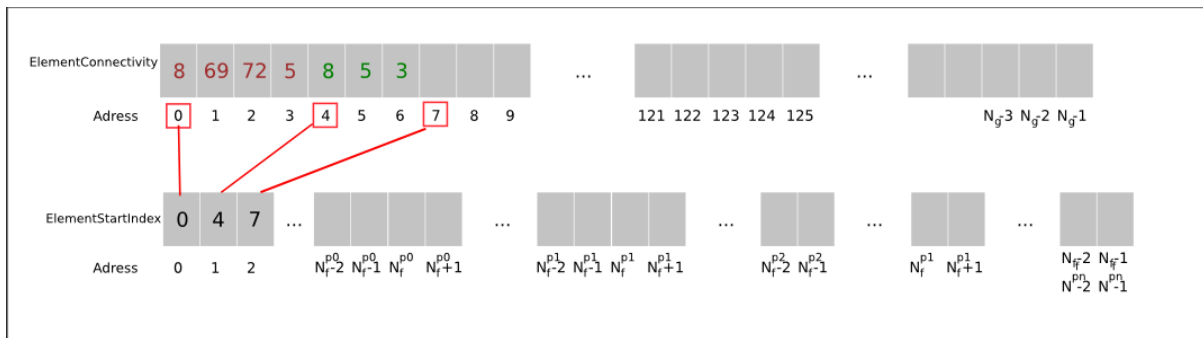


Illustration 5: ElementStartIndex lists the face position in the ElementConnectivity

99 This array lists the position in the ElementConnectivity of the first vertex for each face. This CGNS
 100 node is of type DataArray_t. Its data are consistent and its read can be distributed between N
 101 processors. Its size is N with N being the number of face of the ElementConnectivity face.
 102 The distributed read of this array allows to load partially the ElementConnectivity array.

103 NB: The user should use the ElementConnectivity array size to correctly access the last vertex of
 104 the last face of the ElementConnectivity array.

105 2.1.2 Solution analysis

- 106 1. This solution modifies the CGNS standard.
 107 It strongly modifies the representation of the ElementConnectivity node and it inserts the
 108 new array ElementStartIndex.
 109 These modifications could be challenging for the current CFD database.
- 110 2. Data consistency
 111 The data type of both arrays ElementConnectivity and ElementStartIndex are consistent.
- 112 3. Optmal data size
 113 This solution do not duplicate data thus the global data size is unchanged and stays optimal.

114 4. ElementStartIndex incremental index
115 In this proposal we choose to use the face position in the element connectivity instead of the
116 face number of nodes (ElementStartIndex=[0,4,7,11, ...] instead of [4, 3, 4, ...]).

117 The rationales are all IO read/write oriented and detailed here after:
118 • Processor inter-dependency
119 To access the ElementConnectivity data we need to know the position of the face in the
120 ElementConnectivity array.
121 → “Face vertex number” solution
122 In this configuration we would need to sum over the index array to obtain the face position
123 in the ElementConnectivity array. This means either create a new table or perform the
124 computation as many times as needed.
125 Another problem is that we need the result of the sum of face vertex of processor P_0 to
126 obtain the location of the first vertex of the first face assigned to processor P_1 .
127 This behavior generalizes with the need to sum the face vertex of processors P_0, P_1, \dots, P_{M-1} to
128 obtain the location of the first vertex of the first face assigned to processor P_M .
129 **This is not a complex operation but we do not like the idea of imposing a dependency**
130 **on all previous processors computations.**
131 → ”Face position” solution
132 In this configuration the ElementStartIndex array can be split on P processors and each
133 processor related part directly gives the location of the face vertex from the
134 ElementConnectivity array.
135 • Full vs partial load
136 As stated above we need to know the position of the face in the ElementConnectivity array.
137 → “Face vertex number” solution
138 If we want to access data in processor P_M , then we need to access data on all processors
139 P_0, P_1, \dots, P_{M-1} . This means that we need to load the index array on every processor of inferior
140 rank.
141 → “Face position” solution
142 To access the ElementConnectivity we need the boundary of locations for processor P_M .
143 These boundaries can be partially loaded from the ElementStartIndex array by getting the
144 first element of processor P_M part and the first element of processor P_{M+1} part. These two
145 integers indicate the section of ElementConnectivity to be loaded on proc P_M (1st element of
146 P_M till 1st element of $P_{M+1}-1$)
147 **With this proposal, we need to load only two integers from the ElementStartIndex**
148 **array on a processor to fully access the ElementConnectivity array.** Partial load for
149 CGNS/HDF5 is available in CHLone for example.
150 • direct access to a face connectivity
151 For some specific unstructured software, it is interesting to have a simple access to a specific
152 face in the ElementConnectivity array.
153 This is only allowed with the “Face Position” solution.

154 5. ElementStartIndex content
155 Using an incremental index -the face position in the ElementConnectivity array- could lead
156 to large numbers in the ElementStartIndex array. However, these numbers are limited to the
157 size of the ElementConnectivity array as they give access to addresses of that array. We will
158 probably reach the ElementConnectivity array limits first.

159 **3 Example of extension**




Illustration 6: Simple configuration demonstrating the NGON proposal.

Node Name	Dimensions	Bytes
ElementConnectivity	740	2960
ElementConnectivity	555	2220

Illustration 7: Previous vs new ElementConnectivity array

Node Name	Dimensions	Bytes
ElementStartIndex	185	740

Illustration 8: ElementStartIndex array

160 **4 Conclusions**

161 This modification proposal for the NGON representation addresses the HPC issue due to the
162 interlaced data representation of the unstructured connectivity description. The above solution
163 optimizes the data representation for parallel IO and has a low impact on the SIDS to un-interlace
164 data.

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

168 **5 Annexe : Document modification list**

- 169 1. Following Marc Poinot remarks:
 - 170 • Rename array from FaceConnectivityPosition to ElementStartIndex (text and figures)
 - 171 • Reformulate a sentence describing the ElementStartIndex at line 101
 - 172 • Remove remark concerning array size limitation in section 2.1.2.5