EARN Document

Title:	About GRIBM and PAROUTE
Author:	D. Bovio
Date:	1990/09/06
Committee:	EARN Routing Project Group
Document:	RPG1 90
Revision number:	2
Supersedes:	1
Status:	Final
Maintainer:	D. Bovio
Access:	Public

PREFACE

This is a report about some routing tests done with GRIBM (in the text also GR) and PAROUTE (in the text also PA).
The aim of the test was to verify the differences between these products and if there exist any incompatibility (creation of loops) with the use of both in the same network. For this purpose a set of dummy networks as been used as well as the real network for few cases (BITEARN NODES file 9008).
This report does not draw any conclusion about which product is better or if only one (and which one) should be used in the future (the format of the output tables has not been considered anyway but should be taken in account). This topic is matter of discussion for the RPG list and the RTGs authors.

TESTS WITH DUMMY NETWORKS

A set of 6 dummy network has been used (in the text referred as DNET10A, DNET10B, DNET100, DNET101, DNET102, DNET103) with different topology. Later on also 4 variations of the DNET102 (DNET102A/B/C/D) have been used in order to better look at the tie break logic.
All set has been defined with equal linkspeed for all links (9L). The input tags used are: link, linkspeed (blocked 9L) and nodenumber. The nodenumber has been added because PA utilize it in order to break ties. (GR uses only link and linkspeed and does not use the nodenumber at all, the tables produced by GR in fact are the same both if the nodenumber is defined or not).

Networks DNET10A and DNET10B

For both networks are shown the topology the nodenumbers and the routing matrix produced by GR and PA.

==

```
NETWORK: DNET10A              Nodenumbers:  A  0001
                                            B  0002
                                            X  0003
       X                                    Y  0004
      . .
     .   .
    .     .
   .       .
```

```
A         B
 .       .
  .     .
   .   .
    . .
     Y
```

All links = 9L

GRIBM

FROM

TO	A	B	X	Y
A	-	-Y-	*	*
B	-Y-	-	*	*
X	*	*	-	A
Y	*	*	A	-

PAROUTE

FROM

TO	A	B	X	Y
A	-	-X-	*	*
B	-X-	-	*	*
X	*	*	-	A
Y	*	*	A	-

-N- = Asymmetry between the two RTGs

===

NETWORK: DNET10B

Nodenumbers: A 0001
B 0002
X 0003
Y 0004
Z 0005

```
     X
    . .
   .   .
  .     .
 .       .
A....Z....B
 .       .
  .     .
   .   .
    . .
     Y
```

All links = 9L

GRIBM

FROM

TO	A	X	Y	Z	B
A	-	*	*	*	-Z-
X	*	-	A	A	*
Y	*	A	-	A	*
Z	*	A	A	-	*
B	-Z-	*	*	*	-

PAROUTE

FROM

TO	A	X	Y	Z	B
A	-	*	*	*	-X-
X	*	-	A	A	*
Y	*	A	-	A	*
Z	*	A	A	-	*
B	-X-	*	*	*	-

-N- = Asymmetry between the two RTGs

==

These two simple networks just show that the two RTGs have a different way to break a tie. Apparently, anyway, among different paths (i.e. Z-A-X and Z-B-X in DNET10B) somehow both take the same decision. Although this does not mean that the logic is similar as shown by the further tests.

Network DNET100/101/102/103

For each network are shown the topology, the nodenumbers, the routing matrix produced by GR and PA and a table showing the traffic distribution in the network considering that each node sends one file to all the others (for a total amount of 42 files (6*7) sent over the net). For each link the number of file transmitted is quoted.

==

NETWORK: DNET100

Nodenumbers:

A	0001
B	0002
C	0003
D	0004
E	0005
F	0006
G	0007

All links = 9L

```
      A.........B
     . .       . .
    .   .     .   .
   .     .   .     .
  .       . .       .
  F........G........C
  .       . .       .
   .     .   .     .
    .   .     .   .
     . .       . .
      E.........D
```

GRIBM

```
     FROM
     | A | B | C | D | E | F | G |
TO --+---+---+---+---+---+---+---+
   A | - | * |-G-| G |-G-| * | * |
   --+---+---+---+---+---+---+---+
   B | * | - | * |+G+| G |-G-| * |
   --+---+---+---+---+---+---+---+
   C |-G-| * | - | * | D | G | * |
   --+---+---+---+---+---+---+---+
   D | G ||C|| * | - | * |+G+| * |
   --+---+---+---+---+---+---+---+
   E |-G-| G | D | * | - | * | * |
   --+---+---+---+---+---+---+---+
   F | * |-G-| G ||E|| * | - | * |
   --+---+---+---+---+---+---+---+
   G | * | * | * | * | * | * | - |
   --+---+---+---+---+---+---+---+
```

PAROUTE

```
     FROM
     | A | B | C | D | E | F | G |
TO --+---+---+---+---+---+---+---+
   A | - | * |-B-| G |-F-| * | * |
   --+---+---+---+---+---+---+---+
   B | * | - | * |-C-| G |-A-| * |
   --+---+---+---+---+---+---+---+
   C |-B-| * | - | * | D | G | * |
   --+---+---+---+---+---+---+---+
   D | G | C | * | - | * |-E-| * |
   --+---+---+---+---+---+---+---+
   E |-F-| G | D | * | - | * | * |
   --+---+---+---+---+---+---+---+
   F | * |-A-| G | E | * | - | * |
   --+---+---+---+---+---+---+---+
   G | * | * | * | * | * | * | - |
   --+---+---+---+---+---+---+---+
```

```
 *   = Link
|N|  = Asymmetry within the RTG itself
-N-  = Asymmetry between the two RTGs
```

+N+ = Both

Traffic distribution

```
9 |  |  |  |  |  |  |  |  |  |  |  |  |
8 |  |  |  |  |  |  |--|  |  |  |  |  |
7 |  |  |  |  |  |  |  |--|  |  |  |--|  - GR
6 |++|++|++|++|++|++|  |  |--|--|--|  |
5 |  |  |--|--|  |  |  |  |  |  |  |  |
4 |  |  |  |  |  |  |++|++|++|++|++|++|  + PA
3 |  |--|  |  |--|  |  |  |  |  |  |  |
2 |--|  |  |  |  |--|  |  |  |  |  |  |
1 |  |  |  |  |  |  |  |  |  |  |  |  |
  +--+--+--+--+--+--+--+--+--+--+--+--+
   AB BC CD DE EF FA GA GB GC GD GE GF
```

===

NETWORK: DNET101

```
      A.........B
     .         . .
    .         .   .
   .         .     .
  .         .       .
  F        G........C
  .         . .     .
   .       .   .   .
    .     .     . .
     . .         . .
      E.........D
```

Nodenumbers:

A	0001
B	0002
C	0003
D	0004
E	0005
F	0006
G	0007

All links = 9L

GRIBM

```
     FROM
     | A | B | C | D | E | F | G |
TO --+---+---+---+---+---+---+---+
   A | - | * | B |+G+| F | * | B |
   --+---+---+---+---+---+---+---+
   B | * | - | * |+G+| G | A | * |
   --+---+---+---+---+---+---+---+
   C | B | * | - | * | D | A | * |
   --+---+---+---+---+---+---+---+
   D ||B|||C|| * | - | * | E | * |
   --+---+---+---+---+---+---+---+
   E | F | G | D | * | - | * | * |
   --+---+---+---+---+---+---+---+
   F | * | A | B | E | * | - | E |
   --+---+---+---+---+---+---+---+
   G | B | * | * | * | * | E | - |
   --+---+---+---+---+---+---+---+
```

PAROUTE

```
     FROM
     | A | B | C | D | E | F | G |
TO --+---+---+---+---+---+---+---+
   A | - | * | B |-C-| F | * | B |
   --+---+---+---+---+---+---+---+
   B | * | - | * |-C-| G | A | * |
   --+---+---+---+---+---+---+---+
   C | B | * | - | * | D | A | * |
   --+---+---+---+---+---+---+---+
   D | B | C | * | - | * | E | * |
   --+---+---+---+---+---+---+---+
   E | F | G | D | * | - | * | * |
   --+---+---+---+---+---+---+---+
   F | * | A | B | E | * | - | E |
   --+---+---+---+---+---+---+---+
   G | B | * | * | * | * | E | - |
   --+---+---+---+---+---+---+---+
```

```
 *   = Link
|N|  = Asymmetry within the RTG itself
-N-  = Asymmetry between the two RTGs
+N+  = Both
```

Traffic distribution

```
13 |  |  |  |  |  |  |  |  |  |  |
12 |oo|  |  |  |  |  |  |  |  |  |
11 |  |  |  |  |  |  |  |  |  |  |
10 |  |++|  |  |  |  |  |  |  |  |
 9 |  |  |  |  |  |  |  |  |  |  |
 8 |  |--|++|  |oo|oo|--|  |  |  |
 7 |  |  |  |  |  |  |  |  |  |  |
 6 |  |  |--|oo|  |  |++|  |  |oo|  - GR
 5 |  |  |  |  |  |  |  |  |  |  |  + PA
 4 |  |  |  |  |  |  |  |  |--|  |  oo Both
 3 |  |  |  |  |  |  |  |  |  |  |
 2 |  |  |  |  |  |  |  |oo|++|  |
 1 |  |  |  |  |  |  |  |  |  |  |
   +--+--+--+--+--+--+--+--+--+--+
    AB BC CD DE EF FA GB GC GD GE
```

NETWORK: DNET102

```
     A.........B
    .         . .
   .         .   .
  .         .     .
 .         .       .
F         G         C
 .         .       .
  .         .     .
   .         .   .
    .         . .
     E.........D
```

Nodenumbers:

A	0001
B	0002
C	0003
D	0004
E	0005
F	0006
G	0007

All links = 9L

GRIBM

<table>
<tr><th>TO \ FROM</th><th>A</th><th>B</th><th>C</th><th>D</th><th>E</th><th>F</th><th>G</th></tr>
<tr><td>A</td><td>-</td><td>*</td><td>B</td><td>+G+</td><td>F</td><td>*</td><td>B</td></tr>
<tr><td>B</td><td>*</td><td>-</td><td>*</td><td>+G+</td><td>-F-</td><td>A</td><td>*</td></tr>
<tr><td>C</td><td>B</td><td>*</td><td>-</td><td>*</td><td>D</td><td>A</td><td>-D-</td></tr>
<tr><td>D</td><td>|B|</td><td>|C|</td><td>*</td><td>-</td><td>*</td><td>E</td><td>*</td></tr>
<tr><td>E</td><td>F</td><td>A</td><td>D</td><td>*</td><td>-</td><td>*</td><td>D</td></tr>
<tr><td>F</td><td>*</td><td>A</td><td>B</td><td>E</td><td>*</td><td>-</td><td>B</td></tr>
<tr><td>G</td><td>B</td><td>*</td><td>-D-</td><td>*</td><td>D</td><td>A</td><td>-</td></tr>
</table>

PAROUTE

<table>
<tr><th>TO \ FROM</th><th>A</th><th>B</th><th>C</th><th>D</th><th>E</th><th>F</th><th>G</th></tr>
<tr><td>A</td><td>-</td><td>*</td><td>B</td><td>-C-</td><td>F</td><td>*</td><td>B</td></tr>
<tr><td>B</td><td>*</td><td>-</td><td>*</td><td>-C-</td><td>+D+</td><td>A</td><td>*</td></tr>
<tr><td>C</td><td>B</td><td>*</td><td>-</td><td>*</td><td>D</td><td>A</td><td>-B-</td></tr>
<tr><td>D</td><td>B</td><td>C</td><td>*</td><td>-</td><td>*</td><td>E</td><td>*</td></tr>
<tr><td>E</td><td>F</td><td>|A|</td><td>D</td><td>*</td><td>-</td><td>*</td><td>D</td></tr>
<tr><td>F</td><td>*</td><td>A</td><td>B</td><td>E</td><td>*</td><td>-</td><td>B</td></tr>
<tr><td>G</td><td>B</td><td>*</td><td>-B-</td><td>*</td><td>D</td><td>A</td><td>-</td></tr>
</table>

* = Link

```
 |N|  = Asymmetry within the RTG itself
 -N-  = Asymmetry between the two RTGs
 +N+  = Both
```

Traffic distribution

```
17 |  |  |  |  |  |  |  |  |
16 |--|  |  |  |  |  |  |  |
15 |++|  |  |  |  |  |  |  |
14 |  |  |  |  |  |  |  |  |
13 |  |++|  |  |  |  |  |  |
12 |  |  |  |  |  |--|  |  |
11 |  |  |  |  |  |++|  |  |
10 |  |  |  |  |  |  |  |  |
 9 |  |  |++|++|  |  |  |  |
 8 |  |--|--|--|--|  |oo|--|  - GR
 7 |  |  |  |  |++|  |  |  |  o Both
 6 |  |  |  |  |  |  |  |  |
 5 |  |  |  |  |  |  |  |  |
 4 |  |  |  |  |  |  |  |++|  + PA
 3 |  |  |  |  |  |  |  |  |
 2 |  |  |  |  |  |  |  |  |
 1 |  |  |  |  |  |  |  |  |
   +--+--+--+--+--+--+--+--+
    AB BC CD DE EF FA GB GD
```

===

NETWORK: DNET103

```
      A.........B
     .         . .
    .         .   .
   .         .     .
  .         .       .
 F         G         C
  .                 .
   .               .
    .             .
     .           .
      E.........D
```

Nodenumbers:

A	0001
B	0002
C	0003
D	0004
E	0005
F	0006
G	0007

All links = 9L

GRIBM

FROM

TO	A	B	C	D	E	F	G
A	-	*	B	C	F	*	B
B	*	-	*	C	-F-	A	*
C	B	*	-	*	D	A	B
D	B	C	*	-	*	E	B
E	F	A	D	*	-	*	B

PAROUTE

FROM

TO	A	B	C	D	E	F	G
A	-	*	B	C	F	*	B
B	*	-	*	C	+D+	A	*
C	B	*	-	*	D	A	B
D	B	C	*	-	*	E	B
E	F	|A|	D	*	-	*	|B|

```
--+---+---+---+---+---+---+---+      --+---+---+---+---+---+---+---+
F | * | A | B | E | * | - | B |      F | * | A | B | E | * | - | B |
--+---+---+---+---+---+---+---+      --+---+---+---+---+---+---+---+
G | B | * | B | C |-F-| A | - |      G | B | * | B | C |+D+| A | - |
--+---+---+---+---+---+---+---+      --+---+---+---+---+---+---+---+

  *   = Link
 |N|  = Asymmetry within the RTG itself
 -N-  = Asymmetry between the two RTGs
 +N+  = Both
```

```
19 |  |  |  |  |  |  |  |              Traffic distribution
18 |--|  |  |  |  |  |  |
17 |  |  |  |  |  |  |  |
16 |++|++|  |  |  |  |  |
15 |  |  |  |  |  |  |  |
14 |  |--|  |  |  |--|  |
13 |  |  |  |  |  |  |  |
12 |  |  |++|  |  |++|oo|  + PA    o Both
11 |  |  |  |  |  |  |  |  - GR
10 |  |  |--|  |--|  |  |
 9 |  |  |  |  |  |  |  |
 8 |  |  |  |++|++|  |  |
 7 |  |  |  |  |  |  |  |
 6 |  |  |  |--|  |  |  |
 5 |  |  |  |  |  |  |  |
 4 |  |  |  |  |  |  |  |
 3 |  |  |  |  |  |  |  |
 2 |  |  |  |  |  |  |  |
 1 |  |  |  |  |  |  |  |
   +--+--+--+--+--+--+--+
    AB BC CD DE EF FA GB
```

===

Network DNET100 shows that GR introduces itself two asymmetric paths (D-G-B/B-C-D and F-G-D/D-E-F). PA does not introduce any but the use of both leads to at least 3 more asymmetry.
In this case PA gives a better distribution of the traffic.

Network DNET101 shows that GR introduces itself one asymmetry (D-G-B/B-C-D), plus a second one induced by the first (A-B-C-D/D-G-B-A). PA does not introduce any asymmetry and this is the only difference between the two tables. The traffic distribution it is obviously very similar.

Network DNET102 shows for GR exactly the same asymmetry of DNET101. PA introduces one asymmetry itself (E-D-C-B/B-A-F-E). This case has been studied with networks 102A/102B/102C/102D and it is discussed in the next chapter. If the RTGs are used together 2 more asymmetry are introduced.
The traffic distribution shows only two large peeks for GR and a more variable distribution for PA.

Network DNET103 shows one asymmetry introduced by PA itself (the same as before), plus a second one induced

(G-B-A-F-E/E-D-C-B-G). GR does not introduces any asymmetry and also in this case this is the only difference between the two tables.
The traffic distribution is quite variable for both but PA gives lower peeks.

Networks DNET102A/B/C/D (about the tie-break logic)

To see the dependence of PA by the nodenumber 4 different versions of DNET102 have been used. The topology of these networks is the same as the speed of the links, only the nodenumber has been changed. After the topology, the 4 routing tables obtained are shown with the nodenumbers besides and the available paths and values in case of asymmetry or equal "nodenumber weight".

==

NETWORK: DNET102 All links = 9L

```
       A.........B
      .         . .
     .         .   .
    .         .     .
   .         .       .
   F         G       C
   .         .       .
    .         .     .
     .         .   .
      .         . .
       E.........D
```

```
   *  = Link
  |N| = Asymmetry within the RTG itself
```

```
      PAROUTE  DNET102A                   Nodenumbers: A  0001
                                                       B  0002
      FROM                                             C  0003
        | A | B | C | D | E | F | G |                  D  0004
   TO --+---+---+---+---+---+---+---+                  E  0005
      A | - | * | B | C | F | * | B |                  F  0006
      --+---+---+---+---+---+---+---+                  G  0007
      B | * | - | * | C ||D|| A | * |
      --+---+---+---+---+---+---+---+     Path available:  Path chosen:
      C | B | * | - | * | D | A | B |     E -F-A- B =  7   <-- from B
      --+---+---+---+---+---+---+---+     E -D-C- B =  7   <-- from E
      D | B | C | * | - | * | E | * |     E -D-G- B = 11
      --+---+---+---+---+---+---+---+
      E | F ||A|| D | * | - | * | D |
      --+---+---+---+---+---+---+---+
      F | * | A | B | E | * | - | B |
      --+---+---+---+---+---+---+---+
      G | B | * | B | * | D | A | - |
      --+---+---+---+---+---+---+---+
```

```
        PAROUTE   DNET102B                   Nodenumbers: A   0007
                                                          B   0006
        FROM                                              C   0005
         | A | B | C | D | E | F | G |                    D   0004
   TO  --+---+---+---+---+---+---+---+                    E   0003
       A | - | * | B | E | F | * | B |                    F   0002
       --+---+---+---+---+---+---+---+                    G   0001
       B | * | - | * | G | D | A | * |
       --+---+---+---+---+---+---+---+  Path available:   Path chosen:
       C | B | * | - | * | D | E | D |  E -F-A- B =  9
       --+---+---+---+---+---+---+---+  E -D-C- B =  9
       D | F | G | * | - | * | E | * |  E -D-G- B =  5    <-- from B and E
       --+---+---+---+---+---+---+---+
       E | F | G | D | * | - | * | D |
       --+---+---+---+---+---+---+---+
       F | * | A | D | E | * | - | D |
       --+---+---+---+---+---+---+---+
       G | B | * | D | * | D | E | - |
       --+---+---+---+---+---+---+---+
```

```
        PAROUTE   DNET102C                   Nodenumbers: A   0007
                                                          B   0006
        FROM                                              C   0005
         | A | B | C | D | E | F | G |                    D   0004
   TO  --+---+---+---+---+---+---+---+                    E   0003
       A | - | * | B | E | F | * | B |                    F   0002
       --+---+---+---+---+---+---+---+                    G   0011
       B | * | - | * | C ||D|| A | * |
       --+---+---+---+---+---+---+---+  Path available:   Path chosen:
       C | B | * | - | * | D | E | D |  E -F-A- B =  9    <-- from B
       --+---+---+---+---+---+---+---+  E -D-C- B =  9    <-- from E
       D | F | C | * | - | * | E | * |  E -D-G- B = 15
       --+---+---+---+---+---+---+---+
       E | F ||A|| D | * | - | * | D |
       --+---+---+---+---+---+---+---+
       F | * | A | D | E | * | - | D |
       --+---+---+---+---+---+---+---+
       G | B | * | D | * | D | E | - |
       --+---+---+---+---+---+---+---+
```

```
        PAROUTE   DNET102D                   Nodenumbers: A   0001
                                                          B   0005
        FROM                                              C   0003
         | A | B | C | D | E | F | G |                    D   0004
   TO  --+---+---+---+---+---+---+---+                    E   0002
       A | - | * | B | E | F | * | B |                    F   0006
       --+---+---+---+---+---+---+---+                    G   0007
       B | * | - | * | C ||D|| A | * |
       --+---+---+---+---+---+---+---+  Path available:   Path chosen:
       C | B | * | - | * | D | E | D |  E -F-A- B =  7    <-- from B
       --+---+---+---+---+---+---+---+  E -D-C- B =  7    <-- from E
       D | F | C | * | - | * | E | * |  E -D-G- B = 11
       --+---+---+---+---+---+---+---+
       E | F ||A|| D | * | - | * | D |  A -B-C- D =  8
       --+---+---+---+---+---+---+---+  A -F-E- D =  8    <-- from A and D
       F | * | A | D | E | * | - | D |  A -B-G- D = 12
       --+---+---+---+---+---+---+---+
       G | B | * | D | * | D | E | - |  F -E-D- C =  6    <-- from F and C
       --+---+---+---+---+---+---+---+  F -A-B- C =  6
```

==

Network DNET102 version A shows one asymmetry as we said before. If we consider the "nodenumber weight" (not taking in account the origin and the destination value) the three possible paths score 7, 7, 11. Apparently PA breaks the tie choosing the lower value but gets in trouble because of the double 7.

Network DNET102B has the nodenumbers sequence inverted. First of all the table generated is very different from DNET102A (10 different routes). The asymmetry is no more there and is easy to see why: the same three paths from E to B now score 9, 9, 5 so the lower one is unique.

Network DNET102C has the same sequence of value than DNET102B excluding the value of the node G that is raised at 11. The table produced in fact is not very different from DNET102B but the asymmetry is there again because the available paths (from E to B) score 9, 9, 15.

Network DNET102D has been built with the aim to introduce as much ties as possible. The table produced is exactly the same produced for DNET102C. The output, anyway, given the entry values (the sequence of the nodenumbers is completely different), should be read in a different way.
The asymmetry from E to B is the same as before (7, 7, 11) but on the other hand the paths from A to D, (8, 8, 12) very similar, do not introduce any asymmetry. Finally also the tie between the paths available from F to C (score 6, 6) it is correctly broken.

Drawing some conclusion about the tie break methods.... PA uses an algorithm based on the nodenumber value (empiric but unique within the data base). The result of the tests seems to show that the logic used is to choose the lower value among the totals obtained for the various available paths (n.b. this is only a guess). In some cases, anyway, this algorithm seems not to work properly, introducing some asymmetry (see DNET102/103). Moreover, having the choice of the best path in case of ties bound to an empiric value, PA can also lead to a different "random" output if e.g. someday an old node dies and a new one with an higher nodenumber is introduced at the same place in the topology (see network DNET102C).
On the other hand GR does not uses the nodenumber to break the ties. The results for all these different networks in fact are always the same. Instead of trying to discover the logic GR uses in this case I have asked directly the author. The answer is that GR does not use any "a priori" logic but the choice of the best path in case of a tie is more or less casual. It depends on the moment the program is making the evaluation and may change time by time if somehow, in another side of the network, a node has been added or deleted. As shown by the results also the GR output is affected sometimes by asymmetric paths.

TESTS WITH BITEARN NODES

Few general purpose tests have been performed using the BITEARN NODES file. Here is the results about speed and CPU usage. The test is about respectively the first 1700 (speed) and 49 (CPU) nodes.
The values in brackets are the average per node.

TEST	GR	PA
SPEED: for 1700 real nodes	46.416s (27.3)	41.919s (24.6)
CPU: use of CPU for 49 nodes	VTIME 8:01s (9.8) TTIME 8:06s	VTIME 6:45s (8.2) TTIME 6:52s

Tables for the nodes PUNFSV2, CUNYVMV2, CEARNV2 and DEARN have been produced. For nodes CUNYVMV2 and CEARNV2 the routing tables produced by GR and PA are exactly the same.
For node PUNFSV2 there are 26 different statement, always for the same square (USCVM/UCBCMSA). Here is the log of the compare:

```
Ready; T=0.01/0.01 16:03:29
DMSCMP179I Comparing PUNFSV2 NET-PA A with PUNFSV2 NET-GR A.
ROUTE CSMCMVAX UCBCMSA
ROUTE CSMCMVAX USCVM
ROUTE HARBORUC UCBCMSA
ROUTE HARBORUC USCVM
ROUTE HARBOR1  UCBCMSA
ROUTE HARBOR1  USCVM
ROUTE HARBOR2  UCBCMSA
ROUTE HARBOR2  USCVM
ROUTE HARBOR3  UCBCMSA
ROUTE HARBOR3  USCVM
ROUTE PEPPCDRM UCBCMSA
ROUTE PEPPCDRM USCVM
ROUTE PEPVAX   UCBCMSA
ROUTE PEPVAX   USCVM
ROUTE UCLAAIS  UCBCMSA
ROUTE UCLAAIS  USCVM
ROUTE UCLABIO  UCBCMSA
ROUTE UCLABIO  USCVM
ROUTE UCLACAD  UCBCMSA
ROUTE UCLACAD  USCVM
ROUTE UCLACH   UCBCMSA
ROUTE UCLACH   USCVM
ROUTE UCLACN1  UCBCMSA
ROUTE UCLACN1  USCVM
ROUTE UCLAGSM  UCBCMSA
ROUTE UCLAGSM  USCVM
ROUTE UCLAIEPI UCBCMSA
ROUTE UCLAIEPI USCVM
ROUTE UCLAMVS  UCBCMSA
ROUTE UCLAMVS  USCVM
ROUTE UCLAMVSX UCBCMSA
ROUTE UCLAMVSX USCVM
ROUTE UCLAPH   UCBCMSA
ROUTE UCLAPH   USCVM
```

```
ROUTE UCLASAUP UCBCMSA
ROUTE UCLASAUP USCVM
ROUTE UCLASP   UCBCMSA
ROUTE UCLASP   USCVM
ROUTE UCLASS   UCBCMSA
ROUTE UCLASS   USCVM
ROUTE UCLASSCF UCBCMSA
ROUTE UCLASSCF USCVM
ROUTE UCLASTRO UCBCMSA
ROUTE UCLASTRO USCVM
ROUTE UCLATMOS UCBCMSA
ROUTE UCLATMOS USCVM
ROUTE UCLAUE   UCBCMSA
ROUTE UCLAUE   USCVM
ROUTE UCLAVMXA UCBCMSA
ROUTE UCLAVMXA USCVM
ROUTE UFRJ     UCBCMSA
ROUTE UFRJ     USCVM
DMSCMP209W Files do not compare.
Ready; T=0.16/0.20 16:04:05
SP CONS CLOSE
```

The situation is as follows:

```
            9L....USCVM.....9L
           .         .         .
          .          .          .
   PUNFSV2          9L           UCLACN1
          .          .          .
           .         .         .
            9L...UCBCMSA....9L
```

PA takes the lower route (UCBCMSA) and GR takes the upper (USCVM). This case is just a "real life" example of what has been showed by the dummy networks.
A test made by U. Giese for this case shows that swapping the :adjnodes. tags for USCVM/UCBCMSA the GR output changes and comes out exactly like the PA one. This requires further tests to be done considering the :adjnodes tag too.

The tables for node DEARN show 49 different statement for two different pairs (DLRMVS/DLRVM and DDAGSI1/DHDURZ1):

```
Ready; T=0.01/0.01 13:18:15
FILEL
DMSCMP179I Comparing DEARN NET-PA A with DEARN NET-GR A.
ROUTE DAAFHT1  DDAGSI1
ROUTE DAAFHT1  DHDURZ1
ROUTE DERDBS5  DLRMVS
ROUTE DERDBS5  DLRVM
ROUTE DERRZE0  DLRMVS
ROUTE DERRZE0  DLRVM
ROUTE DFVLROP1 DLRMVS
ROUTE DFVLROP1 DLRVM
ROUTE DGABLG5P DLRMVS
ROUTE DGABLG5P DLRVM
ROUTE DGACHEM5 DLRMVS
ROUTE DGACHEM5 DLRVM
ROUTE DGAESO51 DLRMVS
ROUTE DGAESO51 DLRVM
ROUTE DGAGRS2A DLRMVS
```

```
ROUTE DGAGRS2A DLRVM
ROUTE DGAIPP1S DLRMVS
ROUTE DGAIPP1S DLRVM
ROUTE DGAIPP5D DLRMVS
ROUTE DGAIPP5D DLRVM
ROUTE DGAIPP5N DLRMVS
ROUTE DGAIPP5N DLRVM
ROUTE DGAMPE5D DLRMVS
ROUTE DGAMPE5D DLRVM
ROUTE DGATUM5P DLRMVS
ROUTE DGATUM5P DLRVM
ROUTE DHDDES5  DDAGSI1
ROUTE DHDDES5  DHDURZ1
ROUTE DHDDKFZ1 DDAGSI1
ROUTE DHDDKFZ1 DHDURZ1
ROUTE DHDEMBL  DDAGSI1
ROUTE DHDEMBL  DHDURZ1
ROUTE DHDEMBL5 DDAGSI1
ROUTE DHDEMBL5 DHDURZ1
ROUTE DHDIBM1  DDAGSI1
ROUTE DHDIBM1  DHDURZ1
ROUTE DHDIBM1A DDAGSI1
ROUTE DHDIBM1A DHDURZ1
ROUTE DHDIHEP1 DDAGSI1
ROUTE DHDIHEP1 DHDURZ1
ROUTE DHDIHEP5 DDAGSI1
ROUTE DHDIHEP5 DHDURZ1
ROUTE DHDMPI5  DDAGSI1
ROUTE DHDMPI5  DHDURZ1
ROUTE DHDMPI5D DDAGSI1
ROUTE DHDMPI5D DHDURZ1
ROUTE DHDMPI5H DDAGSI1
ROUTE DHDMPI5H DHDURZ1
ROUTE DHDMPI5V DDAGSI1
ROUTE DHDMPI5V DHDURZ1
ROUTE DHDMPI50 DDAGSI1
ROUTE DHDMPI50 DHDURZ1
ROUTE DHDMPI52 DDAGSI1
ROUTE DHDMPI52 DHDURZ1
ROUTE DHDPHY5  DDAGSI1
ROUTE DHDPHY5  DHDURZ1
ROUTE DHDSPRI6 DDAGSI1
ROUTE DHDSPRI6 DHDURZ1
ROUTE DHDTRN1  DDAGSI1
ROUTE DHDTRN1  DHDURZ1
ROUTE DHDUB1   DDAGSI1
ROUTE DHDUB1   DHDURZ1
ROUTE DHDURZ2  DDAGSI1
ROUTE DHDURZ2  DHDURZ1
ROUTE DHNFHS1  DDAGSI1
ROUTE DHNFHS1  DHDURZ1
ROUTE DK0GRS2I DLRMVS
ROUTE DK0GRS2I DLRVM
ROUTE DMAFHT1  DDAGSI1
ROUTE DMAFHT1  DHDURZ1
ROUTE DMARUM8  DDAGSI1
ROUTE DMARUM8  DHDURZ1
ROUTE DMZMPI5P DLRMVS
ROUTE DMZMPI5P DLRVM
ROUTE DM0GSF11 DLRMVS
ROUTE DM0GSF11 DLRVM
```

```
ROUTE DM0GSF51 DLRMVS
ROUTE DM0GSF51 DLRVM
ROUTE DM0LRZ01 DLRMVS
ROUTE DM0LRZ01 DLRVM
ROUTE DM0MPB51 DLRMVS
ROUTE DM0MPB51 DLRVM
ROUTE DM0MPF11 DLRMVS
ROUTE DM0MPF11 DLRVM
ROUTE DM0MPI11 DLRMVS
ROUTE DM0MPI11 DLRVM
ROUTE DM0MPI12 DLRMVS
ROUTE DM0MPI12 DLRVM
ROUTE DM0MPI53 DLRMVS
ROUTE DM0MPI53 DLRVM
ROUTE DM0TUI1S DLRMVS
ROUTE DM0TUI1S DLRVM
ROUTE EMBL     DDAGSI1
ROUTE EMBL     DHDURZ1
ROUTE MIPS     DLRMVS
ROUTE MIPS     DLRVM
ROUTE SPRINGER DDAGSI1
ROUTE SPRINGER DHDURZ1
ROUTE TUCHEMIE DLRMVS
ROUTE TUCHEMIE DLRVM
DMSCMP209W Files do not compare.
Ready; T=0.16/0.19 13:19:03
SP CONS CLOSE
```

The situation is as follows (example):

```
 DHDUB1....9L...DHDURZ1...9L          9L....DLRVM.....9L....DM0MPI12
                   .         .       .         .
                   .          .     .          .
                  ZL            DEARN          CC
                   .          .     .          .
                   .         .       .         .
                DDAGSI1...ZL          ZL....DLRMVS
```

For both sides GR chooses the "less-hop" path (upper, 9L+9L). PA takes the lower, also if this means one hop more than the other, considering ZL+ZL+9L and ZL+CC+9L better than 9L+9L respectively.
This difference it bounded to the fact that GR implements this rule (extract from the GR description):

```
>  Attribute LINKSnn:
>
>     The names of adjacent nodes and the resistance value
>     (derived or explicitly given), and the NODERESIST
>     attribute will be used in a shortest path algorithm
>     which is based on the classical Dijkstra algorithm with
>     the modification that direct connections are always
>     taken as best path.  (i.e. a 9.6kb connection will be
>     always preferred over two indirect 64kb lines.)
```

This means that ZL+CC+9L is better than 9L+9L for both RTGs but the "triangle rule" above leads GR to use the 9L+9L path.
Another important thing that should be noted is that the application of the "triangle rule" allows GR to avoid the introduction of one asymmetry in this case. The choice of PA,

in fact, introduces an asymmetry because on the way back (e.g. from DHDUB1 to DEARN) the route taken is obviously trough DHDURZ1 and DLRVM respectively due to the direct link.

CONCLUSION

The results show that the two RTGs give different output when some equal alternative path exist because one of them uses an algorithm (PA) to break ties and the other does not (GR).
Different output are possible also because of the different threshold values and rules used while choosing the best path. In this case PA in particular can introduce some asymmetry (see DEARN ex.).
Although no evidence of incompatibility between the two RTGs (loops) has been founded so far, the results show that both RTGs introduce some asymmetric path itself and due to the differences explained above if used together in the same network they will increase the number of asymmetric paths.

This set of tests is obviously not exhaustive. Comments and suggestions from the RTGs authors and the RPG are welcome. I am also available to perform other different tests if they are suggested and/or considered of interest.