WenQuanYi Micro Hei [Scale=0.9]WenQuanYi Micro Hei Mono song-WenQuanYi Micro Hei sfWenQuanYi Micro Hei "zh" = 0pt plus 1pt

FleetX

åŔŚåÿČ 0.1.0.beta

PaddlePaddle

2021 åźť 02 ælJĹ 22 æŮě

åĹĘåÿČåijŔèő■çżČæęĆè£ř

- æňćè£ŐæĆĺåĚşæşĺéčđæąĺåĹĘåÿČåijŕèő■çżČïijŇæĹŚäžňåÿŇæIJŻèČ¡åÿőåŁľæŕŔäÿĂäÿłçŤĺæĹůèţřäÿŁ

CHAPTER 1

æŢťä¡ŞäżŃçż■äÿŐåĘĚåőźæęĆèğĹ

1.1 æňćè£ŐåĚşæşĺåď˙ğèğĎæĺąæůśåžęå■ęäźăæŁĂæIJŕ

è£Śå■Ąåźť˙æİěïijŇæůśåžęå■ęäźăæŁĂæIJŕäÿ■æŰ■åĹůæŰřèğĘèğL˙ãĂĄèĞłçĎűèŕ■èĺĂãĂĄèŕ■é§şãĂĄæŘ

1.2 éčđæąĺåĹĘåÿČåijŔèő■çżČæŔŘä¿ŻçŽĎæăÿå£ČäžůåĂij

1. æžŘèĞłäžğäÿŽåőđèůţçŽĎçżŔélŇïijŽ
   - éčđæąĺçŽĎåĹĘåÿČåijŔèő■çżČæŁĂæIJŕæžŘèĞłçŹ¿åžęçŽĎäÿŽåŁąåőđèůţïijŇæŸŕçżŔè£ĞèűĚåď˙ğèğĎ
   - éčđæąĺåĹĘåÿČåijŔèő■çżČçżŔè£Ğåőđèůţæč˘ĂélŇçŽĎåžŤçŤĺéćĘå§§åŇĚæŃňèĞłçĎűèŕ■èĺĂåď˙ĎçŘĘïij
2. åőŇåď˙ĞçŽĎåźűèąŇæĺąåijŔïijŽ
   - æŢřæ■őåźűèąŇïijŽéŠĹåŕźäžğäÿŽçŢŇæIJĂåÿÿçŤĺçŽĎæŢřæ■őåźűèąŇæĺąåijŔïijŇéčđæąĺéŠĹåŕźåőđéŹĚ
   - æţĄæřť˙çž£åźűèąŇïijŽéİćåŘŚåijĆæđĎçąňäžűïijŇæţĄæřť˙çž£åźűèąŇèČ¡åď˙§åřĘæĺąåđŃèőąçőŮéČĺåĹĘą
   - æĺąåđŃåźűèąŇïijŽåŕźäžŐèűĚåď˙ğèğĎæĺąåĹĘçśžéŮőécŸïijŇéčđæąĺæŔŘä¿Żèőąçő̊ŮäÿŐå■ŸåĆĺåŘŇæŮ
3. éİćåŘŚäžŚçńŕåIJžæŹŕçŽĎåźűèąŇèő■çżČçżĎäžűïijŽ
   - éčđæąĺéŠĹåŕźéŻĘç¿d˙ç¡ŚçżIJçŐŕåćČãĂĄçąňäžűèő¿åď˙ĞæŕŤè¿Čä¡ŐéĚ■çŽĎåIJžæŹŕæŔŘä¿Żåď˙Žçğ■å
   - éŠĹåŕźäžŚçńŕçőŮåŁŻåĚůæIJL˙åijźæĂğçŽĎçL˙źçĆźïijŇéčđæąĺäź§åğŃçżĹåIJĺæŐćçť˙ćåijźæĂğæůśåžęå■

1.3 åijĂåğŃä¡ăçŽĎåĹĘåÿČåijŔèő■çżČäźŃæŮĚ

- æŢťä¡ŞåĘĚåőźïijŽæĹŚäžňæŐĺè■ŘæĆĺçŻťæŐěæăźæ■őäÿżéąţïijŇæŇL'çĚğçńäèŁĆéąžåžŔéĂŘäÿłæţŔè
- FAQïijŽåŕźäžŐéńŸéćŚåĞžçŐřçŽĎéŮőéćŸïijŇæĹŚäžňäijŽåőŽæIJ§æŢť'çŘEçŻÿåĚşåĘĚåőźåĹřFAQ
- å£ńéĂ§äÿŁæL'ŃïijŽåęĆæđIJæČşæIJĂä¡ŐæĹŘæIJňçŽĎäžĘèğčéčđæąĺçŽĎåĹĘåÿČåijŔèő■çżČïijŇæĹŚ
- GPUåď'ŽæIJžèő■çżČïijŽåęĆæđIJæĆĺåůšçżŔåijĂåğŃä¡£çŤĺGPUè£ŻèąŇåď'ŽæIJžåď'Žå■ąèő■çżČïijŇæ
- åŔĆæŢřæIJ■åŁąåŹĺïijŽä£ąæĄŕæčĂçť'ćãĂĄæŐĺè■ŘçşżçżšécĘå§§åÿÿçŤĺçŽĎåźűèąŇèő■çżČæŰźåijŔïi
- åĚňæIJL'äžŚçŐŕåćČåőđèůţïijŽåęĆæđIJæĆĺåIJĺåĚňæIJL'äžŚäÿŁèůŚèĞłåůšçŽĎGPUåď'Žå■ąäżżåŁąïijŇa
- åijźæĂğèő■çżČïijŽåęĆæđIJåŕźåęĆä¡ŢåĹľçŤĺäžŚçńŕåijźæĂğèţĎæžŘè£ŻèąŇåď'ğèğĎæĺąèŠÿéęŔèő■çżČ
  æIJ■åŁąåđŃåijźæĂğèŠÿéęŔ

1.4 RoadMap

- æĹŚäžňäź§äijŽæŐĺéĂĄåď'ğèğĎæĺąæůšåžęå■ęäźăæŁĂæIJŕécĘå§§æIJĂåL'■æš£çŽĎæŁĂæIJŕåĹřè£ŹéČ
  - è£ŚæIJ§ïijŽå■Čäž£èğĎæĺąæĺąåđŃåŘĆæŢřçŽĎGPUåď'ŽæIJžåď'Žå■ąèő■çżČïijŇæŢňèřůæIJ§å¿Ě

CHAPTER 2

åĚňæIJĽäžŚéĚ■ç¡ő

- TBA(çĞŢæŸŐ)

CHAPTER 3

Kubernetes éČÍç¡š

3.1 æęĆè£ř

åIJÍ kubernetes äÿŁéČÍç¡šåĹĘåÿČåijŔäżżåŁąéIJĂèęĄåőL'èčĚ paddle-operator ãĂĆ paddle-operator éĂŽè£ĞæůżåŁäèĞłåőŽäźL'èţĎæžŘçśżåđŃ (paddlejob) äżěåŔŁéČÍç¡š controller åŠŇäÿĂçşżåĹŮ kubernetes åŐ§çŤ§çżĎäżűçŽĎæŰźåijŔåőđçŐřçőĂå■ŢåőŽäźL'å■şåŔŕè£ŘèąŇ paddle äżżåŁąçŽĎéIJĂæśĆãĂĆ

çŻőåL'■æŤŕæŇĄè£ŘèąŇ ParameterServer (PS) åŠŇ Collective äÿď'çğ■åĹĘåÿČåijŔäżżåŁąïijŇå¡ŞçĎűäź§æŤŕæŇĄè£ŘèąŇå■ŢèŁĆçĆźäżżåŁąãĂĆ

3.2 paddle-operator åőL'èčĚ

3.2.1 åĞĘåď'Ğ

åőL'èčĚ paddle-operator éIJĂèęĄæIJL'åůšçżŔåőL'èčĚçŽĎ kubernetes (v1.16+) éŻEç¿ď'åŠŇ kubectl (v1.16+) åůěåĚůãĂĆ

åŔŕäżěéĂŽè£Ğ *git clone* æĹŰèĂĚåď'■åĹűæŰĞäżűåĘĚåőźä£İå■Ÿäžěäÿŋ æŰĞäżű åĹřæIJňåIJř,

```
deploy
âŤIJâŤĂâŤĂ examples
âŤĆÂăÂă âŤIJâŤĂâŤĂ wide_and_deep.yaml
âŤĆÂăÂă âŤIJâŤĂâŤĂ wide_and_deep_podip.yaml
âŤĆÂăÂă âŤŤâŤĂâŤĂ wide_and_deep_service.yaml
```

(äÿŃéąţçżğçż■)

(çż■äÿŁéąţ)

```
âŤŤâŤĂâŤĂ v1
    âŤIJâŤĂâŤĂ crd.yaml
    âŤŤâŤĂâŤĂ operator.yaml
```

3.2.2 éČÍç¡š CRD

æL'ğèąŇäžěäÿŃåŚ¡äžd'ïijŇ

```
$ kubectl create -f deploy/v1/crd.yaml
```

éĂŽè£ĞäžěäÿŃåŚ¡äžd'æ§ěçIJŃæŸŕåŘęæĹŘåŁ§ïijŇ

```
$ kubectl get crd
NAME                                    CREATED AT
paddlejobs.batch.paddlepaddle.org       2021-02-08T07:43:24Z
```

3.2.3 éČÍç¡š controller åŔŁçŻÿåĚşçżĎäžű

æşÍæĎŔéżŸèőd'éČÍç¡šçŽĎ namespace äÿž paddle-systemïijŇåęĆæđIJåÿŇæIJŻåIJÍèĞłåőŽäźĽ'çŽĎ namespace äÿ■è£ŘèąŇæĹŰèĂĚæŔŘäžd'äżżåŁąïijŇ éIJĂèęĄåĚĹåIJÍ operator.yaml æŰĞäžűäÿ■åŕźåžŤæŻt'æŤź namespace éĚ■ç¡őïijŇåĚűäÿ■

- *namespace: paddle-system* èąÍçd'žèŕěèţĎæžŘéČÍç¡šçŽĎ namespaceïijŇåŔŕçŘĘèğčäÿžçşżçż§ controller namespaceïijŻ
- Deployment èţĎæžŘäÿ■ containers.args äÿ■ *–namespace=paddle-system* èąÍçd'ž controller çŻŚæŐğèţĎæžŘæĽĂåIJÍ namespaceïijŇå■şäżżåŁąæŔŘäžd' namespaceãĂĆ

æL'ğèąŇäžěäÿŃéČÍç¡šåŚ¡äžd'ïijŇ

```
$ kubectl create -f deploy/v1/operator.yaml
```

éĂŽè£ĞäžěäÿŃåŚ¡äžd'æ§ěçIJŃéČÍç¡šçżŞæđIJåŠŇè£ŘèąŇçŁűæĂĄïijŇ

```
$ kubectl -n paddle-system get pods
NAME                                         READY   STATUS    ␣
↪RESTARTS   AGE
paddle-controller-manager-698dd7b855-n65jr   1/1     Running   0   ␣
↪       1m
```

éĂŽè£Ğæ§ěçIJŃ controller æŮěå£ŮäžěçąőäŁİè£ŘèąŇæ■čåÿÿïijŇ

```
$ kubectl -n paddle-system logs paddle-controller-manager-
↪698dd7b855-n65jr
```

æŔŘäžd' demo äżżåŁąæ§ěçIJŃæŢĹæđIJïijŇ

```
$ kubectl -n paddle-system create -f deploy/examples/wide_and_deep.
↪yaml
```

æ§ěçIJŃ paddlejob äżżåŁąçŁűæĂĄ, pdj äÿž paddlejob çŽĎçijľåĘŹïijŇ

```
$ kubectl -n paddle-system get pdj
NAME                        STATUS      MODE    PS     WORKER    AGE
wide-ande-deep-service      Completed   PS      2/2    0/2       4m4s
```

äžěäÿŁä£ąæĄŕåŔŕäžěçIJŃåĞžïijŽèő■çżČäżżåŁąåůšçżŔæ■čçąőåőŇæĹŘïijŇèŕěäżżåŁąäÿž ps æĺąåijŔïijŇéĚ■ç¡őéIJĂæśĆ 2 äÿł pserver, 2 äÿłåIJĺè£ŘèąŇïijŇéIJĂæśĆ 2 äÿł wokerïijŇ0 äÿłåIJĺè£ŘèąŇïijĹåůšåőŇæĹŘéĂĂåĞžïijL˜ãĂĆ åŔŕéĂŽè£Ğ cleanPodPolicy éĚ■ç¡őäżżåŁąåőŇæĹŘ/åď śèt˘ěåŘÕçŽĎ pod åĹäéŹď˘ç■ŰçŢěïijŇèŕęèğĄäżżåŁąéĚ■ç¡őãĆ

æ§ěçIJŃ pod çŁűæĂĄïijŇ

```
$ kubectl -n paddle-system get pods
```

3.2.4 å■ÿè¡¡

éĂŽè£ĞäžěäÿŃåŚ¡äžď˘å■ÿè¡¡éČĺç¡šçŽĎçżĎäžűïijŇ

```
$ kubectl delete -f deploy/v1/crd.yaml -f deploy/v1/operator.yaml
```

æşĺæĎŔïijŽéĞ■æŰřåőL'èčĚæŮűïijŇåžžèőőåĚĹå■ÿè¡¡åĘ■åőL'èčĚ

3.3 paddlejob äżżåŁąæŔŘäžď˘

åIJĺäÿŁè£řåőL'èčĚè£ĞçĺŃäÿ■ïijŇæĹŚäžňä¡£çŤĺäžĘ wide-and-deep çŽĎä¿Ńå■Řä¡IJäÿžæŔŘäžď˘äżżåŁąæijŢçď˘žïijŇæIJňèŁĆèŕęçżĘæŔŔè£řäżżåŁąåřĄèčĚåŠŇæŔŘäžď˘æţĄçĺŃä

3.3.1 äžčçăĄåĞĘåď˘Ğ

çď˘žä¿ŃæžŘçăĄåŔŕåIJĺæ■ď˘èÕůå¿ŮïijŇwide_and_deep ïijŇtrain.py ä¡IJäÿžçĺŃåžŔçŽĎåĚěåŔčçĆźãĂĆ

æIJňçď˘žä¿ŃäijŽåIJĺäżżåŁąéŢIJåČŔäÿ■åŇĚåŘńèő■çżČæŢřæ■őïijŇåőđéŹĚåžŢçŤĺè£ĞçĺŃäÿ■äÿĂèĹñäÿ

- äżżåŁąè£ŘèąŇæŮűïijŇçĺŃåžŔéĂŽè£Ğç¡ŚçżIJæŃL'åŘŰæŢřæ■őåĹřæIJňåIJřè£ŻèąŇèő■çżČïijŇèŕěæČĚ
- äżżåŁąè£ŘèąŇæŮűïijŇçĺŃåžŔèŕżåŘŰæIJňåIJřçŽőå¡Ţè£ŻèąŇèő■çżČïijŇèŕěæČĚå¡ćéIJĂèęĄä¡£çŤĺçŤĺ kubernetes æŤřæŇĄçŽĎæŇĆè¡¡å■ŸåĆĺïijŇäÿĂèĹñäžžèőőä¡£çŤĺ pvc æŁ¡èśąïijŇèŕęçżĘçď˘žä¿ŃèĞĄäÿŃäÿĂåřŔèŁĆãĂĆ

3.3.2 åĹűä¡IJäżżåŁąéŢIJåČŔ

åIJĺ kubernetes äÿ■ä¡£çŤĺéŢIJåČŔéIJĂèęĄæIJL'åŔŕè©£éŮőçŽĎéŢIJåČŔäżŞåžŞïijŇè£ŹéĞŇä¡£çŤĺçŹ¿åž ccr ä¡IJäÿžçđ'žä¿ŃïijŇçŤĺæĹůéIJĂèęĄèĞłåůśéË■ç¡õãĂĆ

çŤĺäžŐçŤ§æĹŘéŢIJåČŔçŽĎ Dockerfile åŠŇäżčçăĄçŽőå¡ŢïijŇ

```
$ ls
Dockerfile   wide_and_deep
```

Dockerfile åĘĚåőźïijŇ

```
$ cat Dockerfile
FROM ubuntu:18.04

RUN apt update && \
    apt install -y python3 python3-dev python3-pip

RUN python3 -m pip install paddlepaddle==2.0.0 -i https://mirror.
↪baidu.com/pypi/simple

## äżěäÿŃæăźæ■őçŤĺæĹůåĘĚåőźä£őæŤź

ADD wide_and_deep /wide_and_deep

WORKDIR /wide_and_deep

ENTRYPOINT ["python3", "train.py"]
```

çŤĺæĹůåŔŕæăźæ■őåőđéŹĚæČĚåĘţæŻť'æŤźåĘĚåőźåŠŇåőL'èčĚécİåđ'Ůä¿İètŰãĂĆ

æşĺæĎŔïijŽä¡£çŤĺ gpu èő■çżČæŮűéIJĂèęĄ

- åőL'èčĚ gpu çL'ĹæIJňçŽĎ paddlepaddle åŠŇçŻÿåĚşçżĎäžűæĹŰéĂL'çŤĺ åőŸæŰź docker ä¡IJäÿžå§žçąĂéŢIJåČŔæĹŰçŐŕåćČ;
- éIJĂèęĄåIJĺéŽĘç¿đ'äÿ■åőL'èčĚåě¡åŕźåžŤ éľśåŁĺ åŠŇ åůěåĚůåŇĚ æŤŕæŇĄãĂĆ

åĹűä¡IJéŢIJåČŔ

```
docker build -t registry.baidubce.com/kuizhiqing/demo-wide-and-
↪deep:v1 .
```

æŔŘäžđ'éŢIJåČŔ (éIJĂèęĄåĚůæIJL'åŕźåžŤæİČéŹŘ)

```
docker push registry.baidubce.com/kuizhiqing/demo-wide-and-deep:v1
```

3.3.3 éĚ■ç¡őäżżåŁą

åĞĘåđ'ĞéĚ■ç¡õæŰĞäżűïijŇ

```
$ cat pdj.yaml
apiVersion: batch.paddlepaddle.org/v1
kind: PaddleJob
metadata:
  name: wide-ande-deep
spec:
  withGloo: 1
  intranet: PodIP
  cleanPodPolicy: OnCompletion
  worker:
    replicas: 2
    template:
      spec:
        containers:
          - name: paddle
            image: registry.baidubce.com/kuizhiqing/demo-wide-and-
↪deep:v1
  ps:
    replicas: 2
    template:
      spec:
        containers:
          - name: paddle
            image: registry.baidubce.com/kuizhiqing/demo-wide-and-
↪deep:v1
```

èŕť æŸŐïijŽ

- æŔŘäžďåŚ¡åŘ■éIJĂèęĄåŤŕäÿĂïijŇåęĆæđIJå■ŸåIJĺåĘšçŁĄèřůåĚĹåĹăéŹďåŐ§ paddlejob çąőä£İåůšçżŔåĹăéŹďåĘ■æŔŘäžď;
- ps æĺąåijŔæŮűéIJĂèęĄåŘŇæŮűéĚ■ç¡ő ps åŠŇ workerïijŇcollective æĺąåijŔæŮűåŔłéIJĂèęĄéĚ■ç¡ő worker å■şåŔŕïijŻ
- withGloo åŔŕéĂĽ'éĚ■ç¡őäÿž 0 äÿ■åŘŕçŤĺïijŇ 1 åŔłåŘŕåŁĺ worker çńŕïijŇ 2 åŘŕåŁĺåĚĺéČĺ(workerçńŕåŠŇServerçńŕ)ïijŇ åżžèőőèő¿ç¡ő 1ïijŻ
- cleanPodPolicy åŔŕéĂĽ'éĚ■ç¡őäÿž Always/Never/OnFailure/OnCompletionïijŇèąĺçď'žäżżåŁąçżĹæ■ćïi podïijŇèřČèŕŢæŮűåżžèőő NeverïijŇçŤ§äžğæŮűåżžèőő OnCompletionïijŻ
- intranet åŔŕéĂĽ'éĚ■ç¡őäÿž Service/PodIPïijŇèąĺçď'ž pod éŮť'çŽĎéĂŽä£qæŰźåijŔïijŇçŤĺæĹůåŔŕäžěäÿ■éĚ■ç¡ő, éžŸèőď'ä¡£çŤĺ PodIPïijŻ
- ps åŠŇ worker çŽĎåĘĚåőźäÿž podTemplateSpecïijŇçŤĺæĹůåŔŕæăźæ■őéIJĂèęĄéĄtäžŐ kubernetes èğĎèŇČæůžåŁăæŽť'åď'ŽåĘĚåőź, åęĆ GPU çŽĎéĚ■ç¡ő.

æŻť'åď'ŽéĚ■ç¡őçď'žä¿ŃïijŇ

```
apiVersion: batch.paddlepaddle.org/v1
kind: PaddleJob
metadata:
  name: wide-ande-deep
```

(äÿŃéątçżğçż■)

(çż■äÿŁéąţ)

```
spec:
  intranet: Service
  cleanPodPolicy: OnCompletion
  worker:
    replicas: 2
    template:
      spec:
        containers:
          - name: paddle
            image: registry.baidubce.com/kuizhiqing/demo-wide-and-
↪deep:v1
            resources:
              limits:
                nvidia.com/gpu: 1
        nodeSelector:
          accelerator: nvidia-tesla-p100
  ps:
    replicas: 2
    template:
      spec:
        containers:
          - name: paddle
            image: registry.baidubce.com/kuizhiqing/demo-wide-and-
↪deep:v1
            resources:
              limits:
                nvidia.com/gpu: 1
        nodeSelector:
          accelerator: nvidia-tesla-p100
```

ä¡£çŤÍ kubectl æŔŘäžď yaml éĚ■ç¡őæŰĞäžűäžëåĹŻåžžäžżåŁąïijŇ

```
$ kubectl -n paddle-system create -f pdj.yaml
```

3.4 æŢřæ■őå■ŸåĆÍ

åIJÍ kubernentes äÿ■ä¡£çŤÍæŇĆè¡¡å■ŸåĆÍåžžèőőä¡£çŤÍ pv/pvc éĚ■ç¡őïijŇèŕęèğĄ persistent-volumes ãĂĆ

è£ŹéĞŇä¡£çŤÍ nfs äžŚçŻŸä¡IJäÿžå■ŸåĆÍä¡IJäÿžçd'žä¿ŃïijŇéĚ■ç¡őæŰĞäžűåęĆäÿŃïijŇ

```
$ cat pv-pvc.yaml
---
apiVersion: v1
kind: PersistentVolume
metadata:
  name: nfs-pv
```

(äÿŃéąţçżğçż■)

(çż■äÿŁéąţ)

```
spec:
  capacity:
    storage: 10Gi
  volumeMode: Filesystem
  accessModes:
    - ReadWriteOnce
  persistentVolumeReclaimPolicy: Recycle
  storageClassName: slow
  mountOptions:
    - hard
    - nfsvers=4.1
  nfs:
    path: /nas
    server: 10.12.201.xx

---
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: nfs-pvc
spec:
  accessModes:
    - ReadWriteOnce
  volumeMode: Filesystem
  resources:
    requests:
      storage: 10Gi
  storageClassName: slow
  volumeName: nfs-pv
```

ä¡£çŤĺäżěäÿŃåŚ¡äžd'åIJĺ namespace paddle-system äÿ■ åĹŻåžž pvc åŘ■äÿž nfs-pvc çŽĎå■ŸåĆĺåčřæŸŐïijŇåőđéŹĚåijŢçŤĺäÿž 10.12.201.xx äÿŁçŽĎ nfs å■ŸåĆĺãĂĆ

```
$ kubectl -n paddle-system apply -f pv-pvc.yaml
```

æşĺæĎŔ pvc éIJĂèęĄçżŠåőŽ namespace äÿŤåŔłèČ¡åIJĺèřě namespace äÿŃä¡£çŤĺãĂĆ

æŔŘäžd' paddlejob äżżåŁąæŮűïijŇéĚ■ç¡ő volumes åijŢçŤĺäżěä¡£çŤĺářźåžŤå■ŸåĆĺïijŇ

```
apiVersion: batch.paddlepaddle.org/v1
kind: PaddleJob
metadata:
  name: paddlejob-demo-1
spec:
  cleanPolicy: OnCompletion
  worker:
    replicas: 2
    template:
      spec:
```

(äÿŃéąţçżğçż■)

(çż■äÿŁéąţ)

```
        restartPolicy: "Never"
        containers:
          - name: paddle
            image: registry.baidubce.com/kuizhiqing/paddle-ubuntu:2.
↪0.0-18.04
            command: ["bash","-c"]
            args: ["cd /nas/wide_and_deep; python3 train.py"]
            volumeMounts:
            - mountPath: /nas
              name: data
        volumes:
          - name: data
            persistentVolumeClaim:
              claimName: nfs-pvc
  ps:
    replicas: 2
    template:
      spec:
        restartPolicy: "Never"
        containers:
          - name: paddle
            image: registry.baidubce.com/kuizhiqing/paddle-ubuntu:2.
↪0.0-18.04
            command: ["bash","-c"]
            args: ["cd /nas/wide_and_deep; python3 train.py"]
            volumeMounts:
            - mountPath: /nas
              name: data
        volumes:
          - name: data
            persistentVolumeClaim:
              claimName: nfs-pvc
```

èŕěçď žä¿Ńäÿ■ïijŇéŢIJåČŔäžĚæŔŘä¿Żè£ŘèąŇçŐŕåćČïijŇèő■çżČäžčçăĄåŠŇæŢřæ■őåİĞéĂŽè£Ğå■Ÿå

CHAPTER 4

åőL'èčĚPaddle

ä¡£çŤĺéčđæąĺè£ŻèąŇåĹĘåÿČåijŔèő■çżČçŽĎæIJĂářŔåőL'èčĚéŻĘåŘĹářśæŸŕåőL'èčĚPaddleãĂĆäżŐPad
2.0çL'ĹæIJňåijĂåğŃïijŇæĹŚäżňéİćåŘŚäÿ■åŘŇçŤĺæĹůç¿d'ä¡ŞæŔŘä¿Żäÿ■åŘŇçśżåđŃçŽĎåĹĘåÿČåijŔèő■ç

- éİćåŘŚçőŮæşŢåůěçÍŃåÿĹäÿžäÿżçŽĎéńŸçžğAPI **paddle.distributed.fleet**ãĂĆ
- éİćåŘŚåĚůæIJL'åĹĘåÿČåijŔèő■çżČåžŢåśĆåůěçÍŃåijĂåŘŚèČ¡åŁŻçŽĎåůěçÍŃåÿĹæŔŘä¿ŻçŽĎAPI **paddle.distributed**ãĂĆ
- æĆĺåŔłéIJĂèęĄåőL'èčĚPaddleïijŇåřśåŔŕäżěèŐůå¿ŮéčđæąĺåŻćéŸ§åőŸæŰźæŔŘä¿ŻçŽĎæL'ĂæIJL'åĹ

åőL'èčĚCPUçL'ĹæIJň

```
pip install paddlepaddle
```

æĹŰèĂĚåőL'èčĚGPUçL'ĹæIJň

```
pip install paddlepaddle-gpu
```

åĚşäžŐåőL'èčĚPaddleïijŇè£ŹéĞŇ æIJL'æŻť'åőŇåd'ĞçŽĎåőL'èčĚæŇĞå■Ůä¿ŻæĆĺåŔĆèĂČãĂĆ

CHAPTER 5

äijŸåŇŰçőŮæşŢ

- TBA(åď'ľåĄě)

CHAPTER 6

Collectiveèő■çżČ

6.1 å£ńéĂ§åijĂåğŃ

6.1.1 Collectiveèő■çżČå£ńéĂ§åijĂåğŃ

æIJňèŁĆåřĘéĞĞçŤÍCVéćĘå§§éİđåÿÿçżŔåĚÿçŽĎæÍąåđŃResNet50äÿžä¿ŃïijŇäżŃçż■åęĆä¡Ţä¡£çŤÍFlee
APIïijĹpaddle.distributed.fleetïijL'åőŇæĹŘCollectiveèő■çżČäżżåŁąãĂĆæŢřæ■őæŰźéİćæĹŚäżňéĞĞçŤÍPadd

6.1.2 çL'ĹælJňèęĄæśĆ

åIJÍçijŰåĘŹåĹĘåÿČåijŔèő■çżČçÍŃåžŔäźŃåL'■ïijŇçŤÍæĹůéIJĂèęĄçąőä£İåůšçżŔåőL'èčĚpaddlepaddle-
2.0.0-rc-cpuæĹŰpaddlepaddle-2.0.0-rc-gpuåŔŁäżěäÿŁçL'ĹælJňçŽĎéčđæąĺåijĂæžŘæąĘæđűãĂĆ

6.1.3 æŞ■ä¡lJæŰźæşŢ

äÿŐå■ŢæIJžå■Ţå■ąçŽĎæŹőéĂŽæÍąåđŃèő■çżČçŻÿæŕŤïijŇCollectiveèő■çżČçŽĎäżčçăĄäÿżèęĄéIJĂèę

1. åŕijåĚěåĹĘåÿČåijŔèő■çżČéIJĂèęĄçŽĎä¿İètŰåŇĚãĂĆ
2. åĹİåğŃåŇŰFleetçŐŕåćČãĂĆ
3. èő¿ç¡őåĹĘåÿČåijŔèő■çżČéIJĂèęĄçŽĎäijŸåŇŰåŹÍãĂĆ
   äÿŃéİćåřĘéĂŘäÿĂè£ŻèąŇèőšèğčãĂĆ

åŕijåĚěä¿ìètŰ

åŕijåĚěå£ĚèęĄçŽĎä¿ìètŰïijŇä¿ŃåęĆåĹĘåÿČåijŔèő■çżČäÿŞçŤĺçŽĎFleet API(paddle.distributed.fleet)ãĂĆ

```
from paddle.distributed import fleet
```

åĹìåğŃåŇŰfleetçŐŕåćČ

åŇĚæŃňåőŽäźL'çijžçIJĄçŽĎåĹĘåÿČåijŔç■ŰçŢěïijŇçĎűåŘŐéĂŽè£ĞårĘåŔĆæŢřis_collectiveèő¿ç¡őäÿ

```
strategy = fleet.DistributedStrategy()
fleet.init(is_collective=True, strategy=strategy)
```

èő¿ç¡őåĹĘåÿČåijŔèő■çżČä¡£çŤĺçŽĎäijŸåŇŰåŹĺ

ä¡£çŤĺdistributed_optimizerèő¿ç¡őåĹĘåÿČåijŔèő■çżČäijŸåŇŰåŹĺãĂĆ

```
optimizer = fleet.distributed_optimizer(optimizer)
```

åőŇæŢť'äżčçăĄ

train_fleet_static.pyçŽĎåőŇæŢť'èő■çżČäżčçăĄåęĆäÿŃæL'Ăçď'žãĂĆ

```
# -*- coding: UTF-8 -*-
import numpy as np
import argparse
import ast
import paddle
# åŕijåĚěå£ĚèęĄåĹĘåÿČåijŔèő■çżČçŽĎä¿ìètŰåŇĚ
import paddle.distributed.fleet as fleet
# åŕijåĚěæĺąåđŃæŰĞäżű
import resnet_static as resnet
import os

base_lr = 0.1    # å■ęäźăçŐĞ
momentum_rate = 0.9 # åĘšéĞŔ
l2_decay = 1e-4 # æİČéĞ■èąřåĞŔ

epoch = 10  #èő■çżČè£■äżčæňąæŢř
batch_size = 32 #èő■çżČæL'źæňąåd'ğåřŔ
class_dim = 10

# èő¿ç¡őäijŸåŇŰåŹĺ
def optimizer_setting(parameter_list=None):
    optimizer = paddle.optimizer.Momentum(
```

(äÿŃéąţçżğçż■)

(çż■äÿŁéąţ)

```
        learning_rate=base_lr,
        momentum=momentum_rate,
        weight_decay=paddle.regularizer.L2Decay(l2_decay),
        parameters=parameter_list)
    return optimizer
# èő¿ç¡őæŢřæ■őèŕżåŔŰåŹÍ
def get_train_loader(feed_list, place):
    def reader_decorator(reader):
        def __reader__():
            for item in reader():
                img = np.array(item[0]).astype('float32').reshape(3,
↪ 224, 224)
                label = np.array(item[1]).astype('int64').reshape(1)
                yield img, label

        return __reader__
    train_reader = paddle.batch(
            reader_decorator(paddle.dataset.flowers.train(use_
↪xmap=True)),
            batch_size=batch_size,
            drop_last=True)
    train_loader = paddle.io.DataLoader.from_generator(
        capacity=32,
        use_double_buffer=True,
        feed_list=feed_list,
        iterable=True)
    train_loader.set_sample_list_generator(train_reader, place)
    return train_loader
# èő¿ç¡őèő■çżČåĞ¡æŢř
def train_resnet():
    paddle.enable_static() # ä¡£èČ¡éİŹæĂĄåŻ¿åŁŞèČ¡
    paddle.vision.set_image_backend('cv2')

    image = paddle.static.data(name="x", shape=[None, 3, 224, 224],␣
↪dtype='float32')
    label= paddle.static.data(name="y", shape=[None, 1], dtype=
↪'int64')
    # èřČçŤÍResNet50æĺąåđŃ
    model = resnet.ResNet(layers=50)
    out = model.net(input=image, class_dim=class_dim)
    avg_cost = paddle.nn.functional.cross_entropy(input=out,␣
↪label=label)
    acc_top1 = paddle.metric.accuracy(input=out, label=label, k=1)
    acc_top5 = paddle.metric.accuracy(input=out, label=label, k=5)
    # èő¿ç¡őèő■çżČètĎæžŘïijŇæIJňä¿Ńä¡£çŤÍGPUètĎæžŘ
    place = paddle.CUDAPlace(int(os.environ.get('FLAGS_selected_gpus
↪', 0)))

    train_loader = get_train_loader([image, label], place)
```

(äÿŃéąţçżğçż■)

(çż■äÿŁéąţ)

```
    #åĹİåğŃåŇŰFleetçŐŕåćČ
    strategy = fleet.DistributedStrategy()
    fleet.init(is_collective=True, strategy=strategy)
    optimizer = optimizer_setting()

    # éĂŽè£ĞFleet␣
↪APIèŐůåŔŰåĹĘåÿČåijŔäijŸåŇŰåŹĺïijŇåřĘåŔĆæŢřäijäåĚěéčđæąĺçŽĎå$žçąĂäijŸåŇŰåŹĺ
    optimizer = fleet.distributed_optimizer(optimizer)
    optimizer.minimize(avg_cost)

    exe = paddle.static.Executor(place)
    exe.run(paddle.static.default_startup_program())

    epoch = 10
    step = 0
    for eop in range(epoch):
        for batch_id, data in enumerate(train_loader()):
            loss, acc1, acc5 = exe.run(paddle.static.default_main_
↪program(), feed=data, fetch_list=[avg_cost.name, acc_top1.name,␣
↪acc_top5.name])
            if batch_id % 5 == 0:
                print("[Epoch %d, batch %d] loss: %.5f, acc1: %.5f,␣
↪acc5: %.5f" % (eop, batch_id, loss, acc1, acc5))
# åŘŕåŁĺèő■çżČ
if __name__ == '__main__':
    train_resnet()
```

6.1.4 è£ŘèąŇçďžä¿Ń

åĄĞèő¿èęĄè£ŘèąŇ2å■ąçŽĎäżżåŁąïijŇéĆčäźĹåŔłéIJĂåIJĺåŚ¡äżd'èąŇäÿ■æL'ğèąŇ:

```
fleetrun --gpus=0,1 train_fleet_static.py
```

æĆĺåřĘçIJŃåĹřæŸ¿çd'žåęĆäÿŃæŮěå£Ůä£ąæĄŕïijŽ

```
----------  Configuration Arguments ----------
gpus: 0,1
heter_worker_num: None
heter_workers:
http_port: None
ips: 127.0.0.1
log_dir: log
...
------------------------------------------------
WARNING 2021-01-04 17:59:08,725 launch.py:314] Not found distinct␣
↪arguments and compiled with cuda. Default use collective mode
launch train in GPU mode
```

(äÿŃéąţçżğçż■)

(çż■äÿŁéąţ)

```
INFO 2021-01-04 17:59:08,727 launch_utils.py:472] Local start 2␣
↪processes. First process distributed environment info (Only For␣
↪Debug):
   ␣
↪+=======================================================================
   |                        Distributed Envs                       ␣
↪Value                      |
   +-----------------------------------------------------------------
↪-----------------------+
   |                  PADDLE_CURRENT_ENDPOINT                  127.0.
↪0.1:17901               |
   |                      PADDLE_TRAINERS_NUM                       ␣
↪ 2                      |
   |                 PADDLE_TRAINER_ENDPOINTS         127.0.0.
↪1:17901,127.0.0.1:18846       |
   |                      FLAGS_selected_gpus                       ␣
↪ 0                      |
   |                        PADDLE_TRAINER_ID                       ␣
↪ 0                      |
   ␣
↪+=======================================================================


...
W0104 17:59:19.018365 43338 device_context.cc:342] Please NOTE:␣
↪device: 0, GPU Compute Capability: 7.0, Driver API Version: 10.2,␣
↪Runtime API Version: 9.2
W0104 17:59:19.022523 43338 device_context.cc:352] device: 0, cuDNN␣
↪Version: 7.4.
W0104 17:59:23.193490 43338 fuse_all_reduce_op_pass.cc:78] Find all_
↪reduce operators: 161. To make the speed faster, some all_reduce␣
↪ops are fused during training, after fusion, the number of all_
↪reduce ops is 5.
[Epoch 0, batch 0] loss: 0.12432, acc1: 0.00000, acc5: 0.06250
[Epoch 0, batch 5] loss: 1.01921, acc1: 0.00000, acc5: 0.00000
...
```

åőŇæŢť2å■ąçŽĎæŮěå£Ůä£ąæĄŕäź§åŔŕåIJĺ./log/çŻőå¡ŢäÿŃæ§ěçIJŃãĂĆäžĘèğčæŻťåď`Ž`fleetru
åŘŕåŁĺåĹĘåÿČåijŔäżżåŁąãĂĆ

å■ŢæIJžåĚńå■ąèő■çżČåŘŕåŁĺåŚ¡äžd'çśżäijijïijŇåŔłéIJĂæ■čçąőæŇĞåőŽgpusåŔĆæŢřå■şåŔŕïijŇåęĆäÿ

```
fleetrun --gpus 0,1,2,3,4,5,6,7 train_fleet_static.py
```

äżŐå■ŢæIJžåď'Ž`å■ąåĹřåď'ŽæIJžåď'Ž`å■ąèő■çżČïijŇåIJĺäžčçăĄäÿŁäÿ■éIJĂèęĄåĄŽäżżä¡ŢæŤźåŁĺïijŇåŔ

```
fleetrun --ips="xx.xx.xx.xx,yy.yy.yy.yy" --gpus 0,1,2,3,4,5,6,7␣
↪train_fleet_static.py
```

6.2 æĂğèČ¡å§žåĞĘ

èŕůèő£éŮőéčđæąĺ Perf Repo èŐůåŔŰéčđæąĺæĂğèČ¡å§žåĞĘæŢřæ■õãĂĆ

6.3 èő¿èőąçżijè£ř

6.3.1 èČŇæŹŕ

çžţèğĆæůśåžęå■ęäźăçŽĎåŔŚåśŢåŔšïijŇäÿ■éŽ¿åŔŚçŐřïijŇå¿Ĺåď´ŽåěäåŞžæĂğçŽĎåůěä¡IJïijŇåĚűåőđæ

äÿ■çőąæŸŕå■ęæIJŕçŢŇè£ŸæŸŕåůěäÿŽçŢŇïijŇéČ¡äÿĂçŻť´èĞť´åŁŻäžŐæŻť´å£ńéĂ§çŽĎèő■çżČæŻť´åď´ğ

æŐěäÿŇæİěæĹŚäżňářšäżŐæĂğèČ¡äijŸåŇŰåŠŇåď´ğæĺąåđŃèő■çżČäÿď´æŰźéİćåĚěæL´ŃïijŇäżŃçż■åĞă

6.3.2 æĂğèČ¡äijŸåŇŰ

åIJÍéĂŽçŤÍGPUåŔŚåÿČäźŃåŘŐïijŇä¡£çŤÍæŸ¿å■ąèő■çżČçěđçżŔç¡ŚçżIJçŽĎçČ■åžęåijĂåğŃçĹĘçĆÿæ

åIJÍåĹĘåÿČåijŔæIJžåŹÍå■ęäźăäÿ■ïijŇæIJĂåÿÿçŤÍçŽĎåźűèąŇæĺąåijŔæŸŕæŢřæ■őåźűèąŇïijŇå■şæŕŔäÿł
reduceæŞ■ä¡IJïijL´äžěåđđçŐřåĚÍåśĂä£ąæĄŕåĚśäžńãĂĆåŔŕäžěçIJŃåĞžïijŇèőąçőŮåŠŇéĂŽä£ąæŸŕåĹĘåÿČåi

èőąçőŮOPèđ■åŘĹ

åIJÍæůśåžęå■ęäźăæąĘæđűäÿ■ïijŇæIJĂå§žæIJňçŽĎèőąçőŮå■ŢåĚČæŸŕçőŮå■ŘïijĹOperatorïijL´ãĂĆä¿Ń

ä¡ĘæŸŕïijŇéśijåŠŇçĘŁæŐŇäÿ■åŔŕåĚijå¿ŮãĂĆæŃěæIJL´åůÍåď´ğçĄţæť´żæĂğæL´ĂèęĄäžŸåĞžçŽĎäžčä
b); out=add(tmp, c)ãĂĆåIJÍè£ŹæăůçŽĎç¡ŚçżIJäÿ■ïijŇæĹŚäżňäijŽåŘŕåŁÍäÿď´æňąèőąçőŮïijŇåźűåijĂ

éŠĹåŕźè£ŹäÿłæŞ■ä¡IJçŽĎäÿĂçğ■äijŸåŇŰæŰźæşŢæŸŕïijŇæĹŚäżňåijĂåŔŚäÿĂäÿłæŤŕæŇĄäÿL´äÿłè¿Şă
b, c)ïijŇä¿£åŔŕäžěå¿ŮåĹřæIJĂçżĹçŽĎçżŞæđIJãĂĆèŕěæŰźæşŢçŽĎäÿĂäÿłéŹĎåŁăăě¡åď´ĎæŸŕè£ŸèŁĆçI

è£Źçğ■æĂİèůŕářśæŸŕæL´ĂèřŞçŽĎèőąçőŮOPèđ■åŘĹïijĹFusion)ïijŇèŕęçżĘåĘĚåőźèŕůåŔĆèĂČ4.1.1ařŔ

æŐěäÿŇæİěçŽĎåĞääÿłařŔèŁĆäijŽçżŞåŘĹäÿĂäÿłçŤ§åŁĺçŽĎä¿Ńå■ŘæİěéŸŘè£ŕåŘĎçğ■äijŸåŇŰç■Űç

éĂŽä£ąOPèđ■åŘĹ

AliceåŠŇBobæL´ĂåIJĺçŽĎåŻ¡åőűçŤţèŕÍåŔůçĂĄå¿ĹéŢ£ïijŇæL´ĂäžěäżäžňåŔŚçŐřæŕŔåĄŽåőŇäÿĂéĄŞ

è£ŹářśæŸŕéĂŽä£ąOPèđ■åŘĹçŽĎæĂİæČşãĂĆæĹŚäżňç§ěéĄŞæŕŔæňąèğęåŔŚéĂŽä£ąéČ¡äijŽæIJL´äÿĂ

éĂŽä£ąOPèđ■åŘĹçŽĎä¡£çŤÍæŰźæşŢèŕůåŔĆèĂČ4.1.2ařŔèŁĆãĂĆ

èőąçőŮåŠŇéĂŽä£ąéĞ■åŕă

æŇL'çĚğäźŃåL'■çŽĎçžęåőŽïijŇåĄŽéćŸå£ńçŽĎäžžïijĹæŕŤåęĆAliceïijL'æŃĺéĂŽçŢţèŕİåŘŐïijŇèęĄç■L

è£Źåřśæ ŸŕéĂŽä£ąåŠŇèőąçőŮéĞ■åŕăçŽĎæĂİæČşãĂĆCUDAäÿ■æIJL'æţĄïijĹstream[2]ïijL'çŽĎæęĆå£

éĂŽä£ąæŃŞæL'ŚäijŸåŇŰ

çŐřåIJĺåĄŽéćŸçŽĎåŻćéŸ§åčőåď ğäžĘïijŇéŹď äžĘAliceåŠŇBobïijŇåŔĹåŁäåĚěäžĘåĞää¡■æŰřåŘŇå■ę

äÿ■åŘŇçŽĎä£ąæĄŕäžď æ■ćç■ŰçŢĕïijŇåŕźåžŢåĹřåĹĘåÿČåijŔèő■çżČäÿ■ïijŇåřśæ Ÿŕäÿ■åŘŇçŽĎéĂŽä£
allreduce[3],Double binary trees[4]ç■L'åď Žçğ■æŃŞæL'ŚçżŞæđĎãĂĆ

éĂŽä£ąæŃŞæL'ŚäijŸåŇŰçŽĎæŻť åď Žä¡£çŢĺæŰźæşŢïijŇèŕůåŔĆèĂČ4.3åřŔèŁĆãĂĆ

æůśåžęæćŕåžęåŐŃçijľ

åĘ■æňąåŻđåĹřäžĚæIJL'AliceåŠŇBobäÿď äžžåĄŽéćŸå■ęäźäçŽĎåIJžæŹŕæİěãĂĆäżŰäžňåIJĺåĄŽéćŸè£Č

äÿŁè£řæĂİèůŕåřśæ ŸŕæůśåžęæćŕåžęåŐŃçijľïijĹDeep Gradient Compression, DGCïijL'çŽĎäÿżèęĄæĂİæČşãĂĆDGCéĂŽè£ĞåřĘæćŕåžęçĺĂçŰŔåŇŰïijŇåIJĺæŕŔè¡őèő■çżČæŮűåŔłéĂL'æŇ correction)ãĂĄæIJňåIJŕæćŕåžęèčĄåL'ł(local gradient cliping)ãĂĄåŁĺéĞŔåŻăå■ŘéĄőèŤ¡(Momentum factor masking) ç■L'åĞäéąźæŁĂåůğãĂĆèŕęçżĘąĘĚåőźåŔŕäžěåŔĆèĂČ4.4.1åřŔèŁĆãĂĆ

Local SGD

AliceåŠŇBobèğL'å¿Ůæšąå£ĚèęĄæŕŔéĄŞécŸéČ¡æL'ŞçŢţèŕİäžď æţĄç■ŤæąĹïijŇåřśçőŮä¡£çŤĺäžĘåL'■è£

äžŐæ Ÿŕäÿď äžžåŔĹæČşåĹřäžĘäÿĂäÿłèČ¡åď §åĞŔåřŚæL'ŞçŢţèŕİæňąæŢřçŽĎç■ŰçŢĕïijŽäżŰäžňåĘşåőŽ

Local SGDåřśæ ŸŕåšžäžŐè£ŹäÿłæĂİèůŕïijŇæIJĂå§žæIJňçŽĎLocal SGDåśđäžŐäÿŁä¿ŃçŽĎçňňäÿĂäÿłç■ŰçŢĕïijŇçŻť æŐěáćđåď ğåŔĆæŢřåŘŇæ■ěçŽĎéŮť éŽŤæİěåĞŔåřŚéĂŽ SGDåŔĹèą■çŢ§åĞžäžĘpost Local SGDåŠŇAdaptive Local SGDäÿď æň¿âĂIJåŁäåijžçL'ĹâĂİïijŽ

- post Local SGDè ő■çżČçŽĎçňňäÿĂäÿłéŸűæőţä£İæŇĄæŕŔçőŮåĞžäÿĂäÿłåŔĆæŢřçŽĎæćŕåžęïijŇåřśåő
- Adaptive Local SGDçŻÿåŕźäžŐpost Local SGDèĂŇèĺĂïijŇæŻť åŁäçĄţæť žïijŇåőČäijŽåŁĺæĂĄèřČæŢ

èŕęçżĘąĘĚåőźåŔŕäžěåŔĆèĂČ4.4.2åřŔèŁĆãĂĆ

èĞłåŁĺæůůåŘĹçš¿åžę

AliceåŠŇBobæIJL'äÿĂäÿłçL'źæőŁèőŕå£ĘèČ¡åŁŻïijŇåřśæ ŸŕåŔŕäžěæŁŁæČşèąĺè£řçŽĎåĘĚåőźïijŇæŔŘç

éŽŔçİĂç§ěèŕĘçĆźåŠŇécŸçŻőèűŁæİěèűŁåď ŽïijŇAliceåŠŇBobèğL'å¿ŮèĎŚå■ŘåŔŚæšL'ïijŇåŔŕèČ¡èĎS

åIJĺåőđéŹĚåžŢçŢĺäÿ■ïijŇåŕźåžŢè£Źçğ■èőŕå£ĘèČ¡åŁŻçŽĎåřśæ ŸŕåŔŁçš¿åžęïijĹFP16ïijL'çśżåđŃïijŇä¡£ loss scalingåŠŇopéżŚçŹ¡åŘ■å■Ţç■L'ç■ŰçŢĕæİěéĄ£åĚ■ãĂĆ

- Dynamic loss scalingïijŽåIJĺAMPèő■çżČè£ĞçĺŃäÿ■ïijŇäÿžäžĘéĄ£åĚ■çš¿åžęäÿŃæžćïijŇæŕŔèő■çżČäÿ

- opéżŚçŹ¡åŘ■å■ŢïijŽéĂŽè£Ğä¡£çŤĺåď ğéĞŔæĺąåđŃåIJĺäÿ■åŘŇåžŤçŤĺåIJžæŹŕäÿ■åŔ■åď■élŇèŕĄåŘŐ

èŕęçżEąĘĚåőźèŕůåŔĆèĂČ4.5åřŔèŁĆãĂĆ

6.3.3 åŔĆèĂČètĎæŰŹ

[1] https://developer.nvidia.com/blog/scaling-deep-learning-training-nccl/

[2] https://developer.nvidia.com/blog/how-overlap-data-transfers-cuda-cc/

[3] https://github.com/baidu-research/baidu-allreduce

[4] https://developer.nvidia.com/blog/massively-scale-deep-learning-training-nccl-2-4

6.4 æĂğèČ¡äijŸåŇŰ

6.4.1 OPèđ■åŘĹïijĹèőąçőŮïijŇéĂŽä£ąïijĽ

èőąçőŮèđ■åŘĹ

åřEąæĺąåđŃç¡ŚçżIJäÿ■éąžåžŔæĽğèąŇçŽĎåď ŽäÿłOPsè£ŻèąŇèđ■åŘĹèČ¡åď §åĞŔåřŚOP èřČåžęçŽĎåijĂéŤĂïijŇæŔŘå■Ğèő■çżČéĂ§åžęãĂĆçŻőåĽ■Fleet äÿ■æŤŕæŇĄåęĆäÿŃ3çğ■çŽĎOP èđ■åŘĹïijŽ

- fuse_all_optimizer_opsïijŽèąĺæŸŐæŸŕåŘęèđ■åŘĹ(fuse) æŸŕåŘęèđ■åŘĹ optimizer_opïijŇäžĚåŕźéČĺåĹĘ optimizer åŘŕçŤĺïijĹSGDãĂĄAdamåŠŇMomentumïijĽãĂĆ
- fuse_elewise_add_act_opsïijŽèąĺæŸŐæŸŕåŘęèđ■åŘĹ(fuse) elementwise_add_opåŠŇactivation_opãĂĆ
- fuse_bn_act_opsïijŽèąĺæŸŐæŸŕåŘęèđ■åŘĹ(fuse) batch_norm_op åŠŇ activation_opãĂĆ

éĂŽåÿÿä¡£çŤĺè£ŹäžŽç■ŰçŢěéČ¡äijŽä¡£æŢťä¡ŞæĽğèąŇè£ĞçĺŃæŻťå£ńãĂĆ

éĂŽä£ąèđ■åŘĹ

AllReduce èđ■åŘĹézŸèőď æČĚåĘţäÿŃäijŽåřĘåŘŇäÿĂlayeräÿ■åŔĆæŢřçŽĎæćŕåžęçŽĎåď ŽäÿłAllReduceæŞ■ä¡IJåŘĹåźűæĹŘäÿĂäÿłãĂĆ æŕŤåęĆåŕźäžŐ fc äÿ■æIJĽWeightåŠŇBiasäÿď äÿłåŔĆæŢřïijŇæĽŞåijĂèŕěéĂĽéąźäźŃåĽ■ïijŇéIJĂèęĄäÿď æ æŞ■ä¡IJãĂĆè£ŹæăůåŔŕäžěåĞŔåřŚæćŕåžęåŘŇæ■ěæŮűçŽĎéĂŽä£ąèĂŮæŮűãĂĆ

æ■ď åď ŰïijŇäÿžæŤŕæŇĄæŻť åď ğçšŠåžęçŽĎåŔĆæŢřæćŕåžęèđ■åŘĹïijŇFleet æŔŘä¿ŻäžĘäžěäÿŃäÿď äÿłéĂĽéąźïijŇçŤĺæĹůåŔŕäžěåIJĺèő■çżČçĺŃåžŔè£ŘèąŇåĽ■åIJĺDistributedStrategyä

- fuse_grad_size_in_MB: æŇĞåőŽæŕŔäÿłAllReduceæŞ■ä¡IJçŽĎæćŕåžęå■ŮèŁĆæŢřïijŇåęĆèŕěåŔĆæŢřç åĹŹæŕŔæňąAllReduceèřČçŤĺäijäè¿Ş16MBçŽĎæćŕåžęãĂĆ èŕěåŔĆæŢřçŽĎçżŔéłŇåĂijäÿžæĂżéĂŽä£ąéĞŔçŽĎå■ĄåĹĘäźŃäÿĂãĂĆ

- fuse_grad_size_in_TFLOPS: æŇĞåőŽæŕŔæňąAllReduceæŞ■ä¡IJçŽĎæIJĂåď ğåśĆæŢřïijŇå■şåĹřè¿¿èŕ åĹŹæIJĂåď ŽæŕŔ50åśĆåĄŽäÿĂæňą fused AllReduceãĂĆ

  æşĺæĎŔïijŽ AllReduceèđ■åŘĹçŻőåL'■äÿ■æŤŕæŇąsparseåŔĆæŢřæćŕåžęãĂĆ

æŞ■ä¡IJåőđèůţ

```
# èőąçőŮèđ■åŘĹ
build_strategy = paddle.static.BuildStrategy()
build_strategy.fuse_elewise_add_act_ops = True
build_strategy.fuse_bn_act_ops = True
build_strategy.fuse_relu_depthwise_conv = True
build_strategy.fuse_broadcast_ops = True
build_strategy.fuse_all_optimizer_ops = True

strategy = paddle.distributed.fleet.DistributedStrategy()
strategy.build_strategy = build_strategy

# éĂŽä£ąèđ■åŘĹ
strategy.fuse_grad_size_in_MB = 16
strategy._fuse_grad_size_in_TFLOPS = 50
strategy.fuse_all_reduce_ops=True
```

äÿŁè£ŕä¿Ńå■Řå■ŸæŤ¿åIJĺïijŽexample/resnet/train_fleet_static_op_fusion.pyãĂĆ åĄĞèő¿èęĄè£ŘèąŇ2å■ąçŽĎäżżåŁąïijŇéĆčäźĹåŔłéIJĂåIJĺåŚ¡äżď'èąŇäÿ■æL'ğèąŇ:

```
fleetrun --gpus=0,1 train_fleet_static_op_fusion.py
```

æĆĺåřĘçIJŃåĹřæŸ¿çď'žåęĆäÿŃæŮěå£Ůä£ąæĄŕïijŽ

```
-----------  Configuration Arguments -----------
gpus: None
heter_worker_num: None
heter_workers:
http_port: None
ips: 127.0.0.1
log_dir: log
...
------------------------------------------------
WARNING 2021-01-19 14:53:04,943 launch.py:316] Not found distinct␣
↪arguments and compiled with cuda. Default use collective mode
launch train in GPU mode
INFO 2021-01-19 14:53:04,945 launch_utils.py:472] Local start 8␣
↪processes. First process distributed environment info (Only For␣
↪Debug):
   ␣
↪+=======================================================================
    |                       Distributed Envs                              ␣
↪Value                      |
```

(äÿŃéąţçžğçż■)

(çż■äÿŁéąţ)

```
    +----------------------------------------------------------------
↪-----------------------+
    |                  PADDLE_CURRENT_ENDPOINT                  127.0.
↪0.1:28355               |
    |                      PADDLE_TRAINERS_NUM                      ␣
↪ 8                      |
    |                  PADDLE_TRAINER_ENDPOINTS  ... 0.1:33653,127.0.
↪0.1:27766,127.0.0.1:16631|
    |                      FLAGS_selected_gpus                      ␣
↪ 0                      |
    |                        PADDLE_TRAINER_ID                      ␣
↪ 0                      |
   ␣
↪+==============================================================================
...
W0119 14:53:16.871562 68031 device_context.cc:362] Please NOTE:␣
↪device: 0, GPU Compute Capability: 7.0, Driver API Version: 10.2,␣
↪Runtime API Version: 9.2
W0119 14:53:16.875859 68031 device_context.cc:372] device: 0, cuDNN␣
↪Version: 7.4.
W0119 14:53:25.973377 68031 build_strategy.cc:116] Currently, fuse_
↪broadcast_ops only works under Reduce mode.
I0119 14:53:27.382609 68031 graph_pattern_detector.cc:101] --- ␣
↪detected 16 subgraphs
I0119 14:53:27.390769 68031 graph_pattern_detector.cc:101] --- ␣
↪detected 16 subgraphs
W0119 14:53:27.407582 68031 fuse_optimizer_op_pass.cc:207] Find␣
↪momentum operators : 161, and 161 for dense gradients. To make␣
↪the speed faster, those optimization are fused during training.
W0119 14:53:27.436177 68031 fuse_all_reduce_op_pass.cc:79] Find all_
↪reduce operators: 161. To make the speed faster, some all_reduce␣
↪ops are fused during training, after fusion, the number of all_
↪reduce ops is 6.
[Epoch 0, batch 0] loss: 0.15131, acc1: 0.00000, acc5: 0.03125
[Epoch 0, batch 5] loss: 1.15416, acc1: 0.00000, acc5: 0.03125
```

6.4.2 éĂŽä£ąéĞ■åŕă

çőĂäżŃ

PaddleçŽĎéĂŽä£ąè£ŻèąŇéĞ■åŕăïijĹoverlapïijL'ïijŇåŕŕäżěæIJL'æŢĹæŕŘå■ĞéĂŽä£ąæŢĹçŐĞãĂĆ

åŐ§çŘĘäżŃçż■

PaddleçŽĎæŢt'ä¡ŞæąĘæđűçŽőåL'■åŕłæIJL'äÿĂäÿłèőąçőŮæţĄïijŇä¡ĘåŕŕäżěæIJL'åd'ŽäÿłéĂŽä£ąæţĄãĂ
éĞ■åŕăéĂŽä£ąæţĄïijŇåŕŕäżěæIJL'æŢĹåĹl'çŤĺéĂŽä£ąåÿęåő¡ïijŇäżŐèĂŇè¿¿åĹřæŻt'äijŸçŽĎéĂŽä£ąæĂğè

cuda-streams-best-practicesãĂĆ

ä¡£çŤÍæŰźæşŢ

FleetåůšçżŕåőđçŐřéĂŽä£ąæţĄoverlapïijŇåŕłéIJĂèő¿ç¡őéĂŽä£ąåŹÍæŢřéĞŕ nccl_comm_num åŕŕäžěåŁăå£ńGPUäźŇéŮť'çŽĎéĂŽä£ąæŢĹçŐĞïijŇåžžèőőå■ŢæIJžèő¿ç¡őäÿž1ïijŇåd'ŽæIJ

```
strategy = fleet.DistributedStrategy()
strategy.nccl_comm_num = 2
strategy.sync_nccl_allreduce=False
```

äÿŁè£řä¿Ńå■Řå■ŸæŤ¿åIJĺïijŽexample/resnet/train_fleet_static_overlap.pyäÿŇéİćïijŇ åĄĞèő¿èęĄè£ŘèąŇ2å■ąçŽĎäżžåŁąïijŇéĆčäźĹåŕłéIJĂåIJĺåŚ¡äžd'èąŇäÿ■æL'ğèąŇ:

```
fleetrun --gpus=0,1 train_fleet_static_overlap.py
```

æĆĺářĘçIJŃåĹřæŸ¿çd'žåęĆäÿŇæŮěå£Ůä£ąæĄŕïijŽ

```
-----------  Configuration Arguments -----------
gpus: 0,1
heter_worker_num: None
heter_workers:
http_port: None
ips: 127.0.0.1
log_dir: log
...
------------------------------------------------
...
   ⌙
 ↪+=========================================================================
    |                          Distributed Envs                          ⌙
 ↪Value                      |
    +--------------------------------------------------------------------
 ↪----------------------+
    |                  PADDLE_CURRENT_ENDPOINT                  127.0.
 ↪0.1:10097               |
    |                      PADDLE_TRAINERS_NUM                       ⌙
 ↪ 2                      |
    |                 PADDLE_TRAINER_ENDPOINTS         127.0.0.
 ↪1:10097,127.0.0.1:59371       |
    |                      FLAGS_selected_gpus                       ⌙
 ↪ 0                      |
    |                        PADDLE_TRAINER_ID                       ⌙
 ↪ 0                      |
   ⌙
 ↪+=========================================================================
...
W0118 21:44:34.542804 70071 device_context.cc:362] Please NOTE:⌙
 ↪device: 0, GPU Compute Capability: 7.0, Driver API Version: 10.2,⌙
 ↪Runtime API Version: 9.2
```

(äÿŇéąçżğçż■)

(çż■äÿŁéąţ)

```
W0118 21:44:34.547377 70071 device_context.cc:372] device: 0, cuDNN␣
↪Version: 7.4.
W0118 21:44:40.178053 70071 fuse_all_reduce_op_pass.cc:79] Find all_
↪reduce operators: 161. To make the speed faster, some all_reduce␣
↪ops are fused during training, after fusion, the number of all_
↪reduce ops is 5.
[Epoch 0, batch 0] loss: 0.14466, acc1: 0.00000, acc5: 0.03125
[Epoch 0, batch 5] loss: 4.00225, acc1: 0.00000, acc5: 0.03125
...
```

6.4.3 éĂŽä£ąæŃŞæĽŚäijŸåŇŰ

åŐ§çŘĘ

- TBA

æŞ■ä¡lJåőđèůţ

Fleet åőđçŐřäžEåžŢåśĆéĂŽè£ĞæŤźåŔŸéĂŽä£ąæŃŞæĽŚïijŇåőđçŐřåĹĘåśĆ allreduceãĂĆçŤĺæĹůåŔłéIJĂèęĄæŇĞåőŽçŻÿåžŢçŽĎDistributedStrategy() çŽĎåijĂåĔşïijŇářšåŔŕäżěéĂĽ'æŃl'äÿ■åŘŇçŽĎéĂŽä£ąæŃŞæĽŚãĂĆ

```
dist_strategy = fleet.DistributedStrategy()
dist_strategy.use_hierarchical_allreduce = True
dist_strategy.hierarchical_allreduce_inter_nranks = 8
```

äÿŁè£ŕä¿Ńå■Řå■ŸæŤ¿åIJĺïijŽexample/resnet/train_fleet_static_communication_topology.pyãĂĆ åĄĞèő¿èęĄè£ŘèąŇ8å■ąçŽĎäżżåŁąïijŇéĆčäźĹåŔłéIJĂåIJĺåŚ¡äżd'èąŇäÿ■æĽ'ğèąŇ:

```
fleetrun --gpus=0,1,2,3,4,5,6,7 train_fleet_static_communication_
↪topology.py
```

æĆĺåřĘçIJŃåĹřæŸ¿çd'žåęĆäÿŃæŮěå£Ůä£ąæĄŕïijŽ

```
-----------  Configuration Arguments -----------
gpus: None
heter_worker_num: None
heter_workers:
http_port: None
ips: 127.0.0.1
log_dir: log
...
------------------------------------------------
...
INFO 2021-01-19 14:58:43,720 launch_utils.py:472] Local start 8␣
↪processes. First process distributed environment info (Only For␣
↪Debug):
```

(äÿŃéąţçżğçż■)

(çż■äÿŁéąţ)

```
   ⌴
↪+=========================================================================
    |                          Distributed Envs                           ⌴
↪Value                       |
    +-----------------------------------------------------------------
↪-----------------------+
    |                  PADDLE_CURRENT_ENDPOINT                  127.0.
↪0.1:53762               |
    |                      PADDLE_TRAINERS_NUM                          ⌴
↪ 8                      |
    |                  PADDLE_TRAINER_ENDPOINTS  ... 0.1:58938,127.0.
↪0.1:54203,127.0.0.1:44221|
    |                      FLAGS_selected_gpus                          ⌴
↪ 0                      |
    |                        PADDLE_TRAINER_ID                          ⌴
↪ 0                      |
    ⌴
↪+=========================================================================
...
W0119 14:58:52.487838 95116 device_context.cc:362] Please NOTE:⌴
↪device: 0, GPU Compute Capability: 7.0, Driver API Version: 10.2,⌴
↪Runtime API Version: 9.2
W0119 14:58:52.493592 95116 device_context.cc:372] device: 0, cuDNN⌴
↪Version: 7.4.
W0119 14:59:01.665702 95116 fuse_all_reduce_op_pass.cc:79] Find all_
↪reduce operators: 161. To make the speed faster, some all_reduce⌴
↪ops are fused during training, after fusion, the number of all_
↪reduce ops is 5.
[Epoch 0, batch 0] loss: 0.13468, acc1: 0.00000, acc5: 0.06250
[Epoch 0, batch 5] loss: 0.18902, acc1: 0.03125, acc5: 0.03125
```

6.4.4 éĂŽä£ąéćŚçŐĞäijŸåŇŰ

åIJÍç¡ŚçżIJåÿęåő¡è¿Čä¡ŐçŽĎèő■çżČåIJžæŹŕïijĹåęĆïijŽ åĚňæIJL'äžŚäÿŁèő■çżČïijŇèĄŤéĆęèő■çżČïijL'
Fleet ä¡IJäÿžPaddleéĂŽçŤÍçŽĎåĹĘåÿČåijŔèő■çżČAPI åőđçŐřäžĘïijŽ `Deep Gradient Compression` åŠŇ `Local SGD` äÿd'çğ■èő■çżČç■ŰçŢěæÌěéŠĹářźæĂğèğčåĘşè£ŹäÿĂéŮőécŸãĂĆ

DGC äijŸåŇŰä¡ŐéĚ■ç¡ŚçżIJçŽĎåĹĘåÿČåijŔGPUèő■çżČ

DGC çőĂäżŃ

åd'ğèğĎæĺąåĹĘåÿČåijŔèő■çżČéIJĂèęĄè¿ČéńŸçŽĎç¡ŚçżIJåÿęåő¡äžëä¿£è£ŻèąŇæćŕåžęçŽĎèĄŽåŘĹæŻ
Deep Gradient Compression åŔŚçŐřïijŽåĹĘåÿČåijŔSGDäÿ■æIJL'99.9%çŽĎæćŕåžęäžd'æ■ćéČ¡æŸŕåĘŮä¡Źç
åőđçŐřäžĘDGCçŽĎçÍĂçŰŔéĂŽä£ąæŰźåijŔïijŇåŔŕæIJL'æŢĹåIJĺä¡ŐéĚ■ç¡ŚçżIJäÿŃè£ŻèąŇGPUåĹĘåÿČåij
åőđçŐřäžĘ DGC èőžæŰĞäÿ■çŽĎ `éćĎçČ■èő■çżČ (warming up training)`, `åŁĺéĞŔä£őæ■č (Momentum Correction)`,

åśĂéČÍæćŕåžęäŁőåL’ł (local gradient clipping),
åŁÍéĞŔåŽäå■ŘæŐl’èŮŔ (Momentum factor masking) ç■Ľ’ç■ŰçŢĕïijŇ
åŠŇ æ■čåĹŹåŇŰéąźäŁőæ■č (Weight Decay Correction)
éĄŁåĚ■çÍĂçŰŔæćŕåžęéĂŽäŁąèő■çżČåÿęæİěçŽĎæIJĂçżĹæÍąåđŃçš¿åžęæ■§åď śãĂĆ

äÿŃéİćåřĘäżŃçż■ DGC çÍĂçŰŔéĂŽäŁąæŰźåijŔçŽĎéĂĆçŤÍåIJžæŹŕãĂĄèŕŢéłŇæŢĹæđIJãĂĄå§žæIJňå

éĂĆçŤÍåIJžæŹŕ

DGCçÍĂçŰŔéĂŽäŁąåIJÍä¡Őåÿęåő¡éĂŽäŁąçŞűéĆĹæŮűäijŽæIJL’è¿Čåď ğçŽĎæĂğèČ¡æŔŘå■ĞïijŇä¡Ę**åI**

èŕŢéłŇæŢĹæđIJ

- æÍąåđŃïijŽFasterRCNN
- çąňäżűïijŽ P40äÿď æIJžåĹĘåÿČåijŔïijŇærŔåŘŕæIJžåŹÍäÿĂå■ąïijŇTCPç¡ŚçżIJætŃèŕŢãĂĆ
- åŔŰ300-700æ■ěèĂŮæŮű/400stepãĂĆ
- çš¿åžęæŮăæ■§ãĂĆ

DGC åŐ§çŘĘçőĂäżŃ

è£ŹéĞŇåřĘçőĂå■ŢäżŃçż■äżŃçż■Fleet DGC äÿ■çŽĎäÿĂäžŻåŐ§çŘĘåŠŇåŕźåžŤåŔĆæŢřåžŤèŕěåęĆä¡Ţ

æćŕåžęçÍĂçŰŔ

DGCçŽĎå§žæIJňæĂİèůŕæŸŕéĂŽè£ĞåŔłäijäéĂĄéĞ■èęĄæćŕåžęïijŇå■şåŔłåŔŚéĂĄåď ğäžŐçżŹåőŽéŸÍ
æ■ćäÿłèğŠåžęïijŇäżŐçŘĘèőžä¿İæ■őäÿŁæİěçIJŃïijŇåśĂéČÍæćŕåžęçť ŕåŁäç■Ľ’åŘŇäžŐéŽŔæŮűéŮť æŐÍçğżå
sizeïijŇïijĹDGCçŻÿå¡ŞäžŐæŕŔäÿĂäÿłæćŕåžęæIJL’èĞłåůśçŽĎbatch sizeïijL’ãĂĆ

åĄĞèő¿ NæŸŕèő■çżČèŁĆçĆźäÿłæŢř, bäÿžå■Ţå■ąbatch
sizeïijŇåśĂéČÍæćŕåžęçť ŕåŁäåŔŕäžěèćńèőd’äÿžbatch sizeäżŐNbåćđåď ğäÿžNbTïijŇåĚűäÿ■TæŸŕäÿd’æňąæŻ
[1] éćĎçČ■èŕČåŔĆ ^^^^^^^^

åŕźäžŐæ■čåÿÿçŽĎèő■çżČïijŇä¡£çŤÍDGCäÿĂèĹňéIJÀè£ŻèąŇéćĎçČ■èő■çżČïijŇåŘęåĹŹåŔŕèČ¡äijŽæI

```
# 1.␣
↪äżě1252äÿłstepäÿžäÿĂäÿłepochïijŇåL’■2äÿłepochsä¡£çŤÍæ■čåÿÿdenseéĂŽäŁąïijŇäŔŐ3ä
↪9%
strategy.dgc_configs = {
    "rampup_begin_step": 1252*2,
    "rampup_step": 1252*3,
    "sparsity": [0.984375, 0.996, 0.999]
}
# 2. åL’■éİć4äÿłepochséČ¡ä¡£çŤÍdenseéĂŽäŁąïijŇäźŃåŘŐézŸèőd’0.
↪999çÍĂçŰŔåžęè£ŘèąŇ
```

(äÿŃéąţçżğçż■)

(çż■äÿŁéąţ)

```
strategy.dgc_configs = {
    "rampup_begin_step": 1252*4,
    "rampup_step": 1,
    "sparsity": [0.999]
}
```

åŕźäžŐFine-tuningèő■çżČïijŇåŔŕæŮăéIJĂéćĎçČ■èő■çżČïijŇäžŐçňň0äÿłepochçŻťæŐěä¡£çŤÍDGCå■şå

```
# äżŐçňň0æ■ěåijĂåğŃDGCçĺĂçŰŔéĂŽä£ą
strategy.dgc_configs = {
    "rampup_begin_step": 0,
    "rampup_step": 1,
    "sparsity": [0.999]
}
```

åśĂéČÍæćŕåžęçťŕåŁăæŤźè£Ż

æ■čåÿÿæČĚåĘţïijŇçÍĂçŰŔæŻťæŰřäijŽäÿěéĞ■å¡śåŞ■æŤűæŢŻæĂğãĂĆDGCäÿ■éĞĞçŤÍåŁÍéĞŔä£őæ■
Correction)åŠŇåśĂéČÍæćŕåžęèčĄåĞŔ(Local Gradient Clip-
ping)ïijŇ åŁÍéĞŔåŻää■ŘæŐľèŮŔïijŇ æ■čåĹŹåŇŰéąźä£őæ■č
4äÿłç■ŰçŢěæÌěèğčåĘşè£ŹäÿłéŮőécŸãĂĆ

åŁÍéĞŔä£őæ■č

äÿŁæŰĞâĂİåśĂéČÍæćŕåžęçťŕåŁăç■L'åŘŇäžŐéŽŔæŮűéŮťæŐÍçğżåćđåŁăbatch
sizeâĂIJçŽĎæŐÍåŕijæšąæIJL'èĂČèŹŚ Momentumå■ŸåIJÍçŽĎæČĚåĘţãĂĆå¡ŞçÍĂçŰŔåžęå¿ĹénŸæŮűïijŇä¡£
åĚňåijŔäijŽæŸ¿èŠŮéŹ■ä¡ŐæÍąåđŃæĂğèČ¡ïijŇæL'Ăäžěé IJĂèęĄåIJÍåŐ§ågŃåĚňåijŔçŽĎå§žçąĂäÿŁåŕźæćŕå

åŁÍéĞŔä£őæ■čä¡£çŤÍéČÍçť'råŁăéĂ§åžęéąźu_tèĂŇéİđçť'råŁăçIJ§åőđçŽĎæćŕåžę$\nabla_{k,t}$æÌëä£őæŤźMomentu
æŰźçÍŃïijŇä£őæ■čåŘŐçŽĎåŁÍéĞŔæŻťæŰřåĚňåijŔåęĆäÿŃïijŽ

$$u_{k,t} = mu_{k,t-1} + \nabla_{k,t}, \quad v_{k,t} = v_{k,t-1} + u_{k,t}, \quad w_{t+1} = w_t - \eta \sum_{k=1}^{N} \text{sparse}\left(v_{k,t}\right)$$

åśĂéČÍæćŕåžęä£őåL'ł

æćŕåžęä£őåL'łæŸŕéŸšæ■ćæćŕåžęçĹEçĆÿçŽĎåÿÿçŤÍæŰźæşŢãĂĆè£ŹæŰźæşŢçŤśPascanuç■L'äžžåIJÍ20
normsåŠŇåď ğäžŐçżŹåőŽéŸĹåĂijæŮűïijŇåršåŕźæćŕåžęrescaleãĂĆæ■čåÿÿæćŕåžęä£őåL'łåIJÍæćŕåžęèĄŽåŘĹ
ãĂĆ

åŁÍéĞŔåŻäå■ŘæŐľèÙŔ

åŽääÿžæŐÍè£§äžĘè¿ČåřŔæćŕåžęæŻťæŰřæİČéĞ■çŽĎæŮűéŮťïijŇæĽĂäžěäijŽæIJĽæİČéĞ■éŹĹæŮğæ

$$Mask \leftarrow |v_{k,t}| > thr, \quad v_{k,t} \leftarrow v_{k,t} \odot \neg Mask, \quad u_{k,t} \leftarrow u_{k,t} \odot \neg Mask$$

æ■ď'æŐľçăĄåŔŕäžěåĄIJæ■ćåžűè£§æćŕåžęäžğçŤ§çŽĎåŁÍéĞŔïijŇéŸšæ■ćéŹĹæŮğæćŕåžęæŁŁæİČéĞ■åijŢåİ

æ■čåĹŹåŇŰ(Weight Decay)éąźä£őæ■č

çśżäijijåŁÍéĞŔä£őæ■čïijŇDGC äÿ■æĹŚäžňåŘŇæăůéIJĂèęĄåŕźæ■čåĹŹåŇŰéąźè£ŻèąŇä£őæ■čæİèèőľå

åŠŇåŁÍéĞŔä£őæĂİèůŕçŽÿåŘŇïijŇä£őæ■čéIJĂèęĄåIJÍåśĂéČÍæćŕåžęäÿŁæůżåŁăåśĂéČÍWeight DecayãĂĆ

$$\nabla_{k,t} = \nabla_{k,t} + \frac{\lambda}{N} w_t$$

äÿŁè£řç■ŰçŢěåůšçżŔåIJÍFleet æąĘæđűäÿ■åőđçŐřïijŇçŤÍæĹůæŮăéążèő¿ç¡őãĂĆ

DGC å£ńéĂ§åijĂåğŃ

äÿŃæŰĞäžěå■ŢæIJžåĚńå■ąäÿŁèő■çżČResNet50 äÿžä¿Ńå■ŘçőĂå■ŢäżŃçż■ Fleet äÿ■ DGC çŽĎä¡£çŤÍãĂĆ åŽääÿž 8åijă GPU çŽĎéĂŽä£ąéČ¡åIJÍåŘŇäÿĂèŁĆçĆźåĘĚïijŇ äÿĂèĹňæČĚåĘţäÿŃæćŕåžęéĂŽä£ąåźűäÿ■äijŽæĹŘäÿžèő■çżČçŽĎçŞűéĆĹïijŇ è£ŹéĞŇåŘłæŸŕäžěåĚűäÿžä¿Ńå■ŘïijŇäżŃçż■Fleet äÿ■ DGC åŔĆæŢřçŽĎèő¿ç¡őãĂĆ

æşĺæĎŔïijŽ

- çąňäžűçŐŕåćČèęĄæśĆïijŽ DGCçŻőåĽ■åŘłæŤŕæŇĄGPUåď'Žå■ąåŘŁåĹĘåÿČåijŔcollectiveèő■çżČïijŇéIJĂèęĄæIJĽçŽÿåžŢçŽĎcudaãĂĄcuDNNãĂĄncclçŐŕåćČãĂĆ
- PaddleçŐŕåćČèęĄæśĆïijŽ DGCåŘłæŤŕæŇĄGPUïijŇæĽĂäžěéIJĂGPUçĽĹæIJňçŽĎ-PaddleãĂĆ

DGC çŻÿåĚşç■ŰçŢě

è£ŹéĞŇåĄĞèő¿ïijŽ1252äÿłstepäÿžäÿĂäÿłepochïijŇåĽ■2äÿłepochsä¡£çŤÍæ■čåÿÿdenseéĂŽä£ąïijŇåŘŐ

- `rampup_begin_step (int)` ïijŽDGC(åŘńéćĎçČ■èő■çżČ)åijĂåğŃçŽĎ step
- `rampup_step (int)` ïijŽDGCäÿ■éćĎçČ■èő■çżČæŇĄçż■çŽĎ step. åęĆæđIJsparsity æŸŕ [0.75, 0.9375, 0.984375, 0.996, 0.999]ïijŇrampup_step èő¿æĹŘ 100æŮűïijŇåIJÍ 0~19 steps æŮű sparsity=0.75ïijŇåIJÍ 20~39 steps æŮű sparsity=0.9375ïijŇäžěæ■ď'çśżæŐÍãĂĆ
- `sparsity (list[float])` ïijŽçÍĂçŰŔåžę threshold, (1 - current sparsity) % çŽĎ-gradient åřĘäijŽèćń allreduceãĂĆ

```
strategy = fleet.DistributedStrategy()

strategy.dgc = True
strategy.dgc_configs = {
    "rampup_begin_step": 1252*2,
    "rampup_step": 1252*3,
    "sparsity": [0.984375, 0.996, 0.999]
}
```

å§žäžŐResNet50ç¡ŚçżIJçŽĎDGCäżčçăĄïijŽexample/resnet/train_fleet_static_dgc.pyãĂĆ

ä¡£çŤĺLocal SGD äijŸåŇŰä¡Őåÿęåő¡äÿŃåĹĘåÿČåijŕèő■çżČ

çőĂäżŃ

åIJĺä¡£çŤĺ distributed SGDè£ŻèąŇæŢřæ■őåźűèąŇçŽĎåĹĘåÿČåijŕèő■çżČæŮűïijŇåÿÿäijŽéĄĞåĹřäžěäÿ

- åĹĘåÿČåijŕèő■çżČçŽĎåŘđåŘŘäijŽåŔŮåĹřéŻEç¿d'äÿ■éŽŕæIJžæĚćèŁĆçĆźïijĹstraggling nodeïijL'åŠŇéĂŽä£ąåžűè£§çŽĎå¡śåŞ■ãĂĆ
- æŢřæ■őåźűèąŇåĹĘåÿČåijŔåćđåd'ğäžĘèő■çżČåőđéŹĚçŽĎbatch sizeïijŇè£Ğåd'ğçŽĎbatch size äijŽå¡śåŞ■æIJĂçżĹçŽĎèő■çżČçš¿åžęãĂĆ

Local SGDéĂŽè£Ğåžűéţ£èŁĆçĆźéŮt'åŘŇæ■ěçŽĎéŮt'éŽŤ(åśĂéČĺåijĆæ■ěèő■çżČ)æİěåĞŕè¡żæĚćèŁ

åŐ§çŘĘäżŃçż■

Local SGDåĞŕè¡żæĚćèŁĆçĆźçŽĎå¡śåŞ■åŠŇåĞŕářŚéĂŽä£ąéćŚçŐĞïijŇæŕŘå■Ğèő■çżČçŽĎåŘđåŘŘ sizeïijL'çŽĎçš¿åžęæ■§åd'śïijŇ[1] åŠŇ [2] åĹĘåĹńæŕŘåĞžäžĘïijŽpost-Local SGD åŠŇ èĞłéĂĆåžŤæ■ěéţ£ (Adaptive Communication) Local SGD ç■ŰçŢěïijŇæİěåĞŕářŚåŔĆæŢřåŘŇæ■ěécŚçŐĞéŹ■ä¡ŐåÿęæİěçŽĎçš¿åžęæ■§åd'śãĂĆåŘŇæ■ěSGDåŠŇLoca SGDåIJĺéĂŽä£ąåŘŇæ■ěäÿŁçŽĎåůőåijĆåęĆäÿŃåŻ¿æL'Ăçd'žãĂĆ

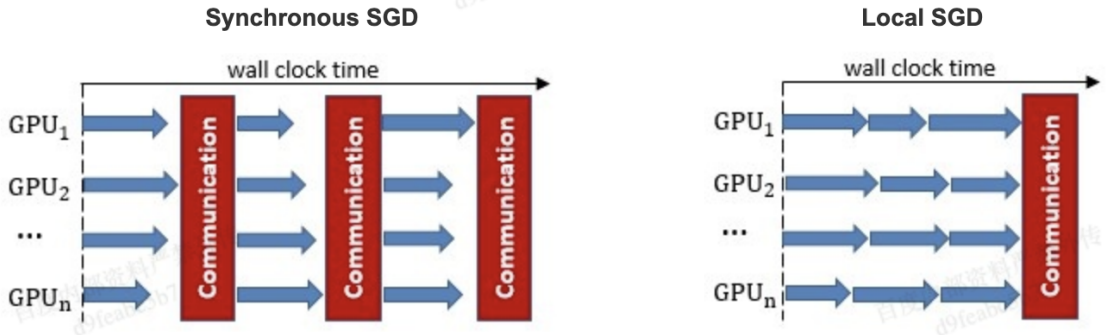

åIJĺLocal SGD èő■çżČäÿ■ïijŇéŻEç¿d'äÿ■çŽĎæŕŘäÿł trainer åŘĎèĞłäijŽçŃňçńŇçŽĎè£ŻèąŇ H äÿłè£đçż■çŽĎ SGD æŽt'æŰřïijŇçĎůåŘŐéŻEç¿d'äÿ■çŽĎæL'ĂæIJL' trainer äijŽè£ŻèąŇéĂŽä£ąïijŇåŘŇæ■ěïijĹaveragingïijL'æL'ĂæIJL' trainers äÿŁçŽĎåŔĆæŢřãĂĆäÿĂäÿłåŔŇ trainersïijŇåŘŇæ■ěéŮt'éŽŹäÿž3

æ■ěéŢ£ïijĹiterationsïijĽçŽĎLocal SGDè£ĞçÍŃåęĆäÿŃåŻ¿æĽĂçď žãĂĆéżĎçż£äÿď æİąèůŕå¿Ďèą[çď žäÿď äÿ
trainers åŘĎèĞłçŽĎ Local SGD æŻť æŰřè£ĞçÍŃïijŇäÿ■éŮť çŽĎèŞİèĽšèůŕå¿Ďèą[çď žåŘŇæ■ěåŘŐçŽĎæĺąå

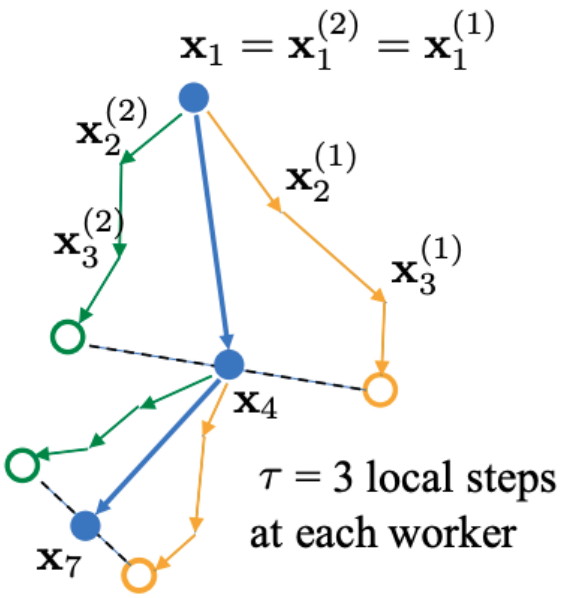

Local SGDäÿ■çŽĎäÿĂäÿłåĚşéŤőéŮőécŸæŸŕåęĆä¡Ţçąőåő ŽåŘĆæŢřåŘŇæ■ěçŽĎéŮť éŽŤ(écŚçŐĞ)ïijŇ

- åćđåď ğåŘĆæŢřåŘŇæ■ěçŽĎéŮť éŽŤåŘŕäžěåĞŘåřŚ trainers éŮť éĂŽä£ąåżűè£§çŽĎå¡śåŞ■æŘŘéńŸèő■çżĈåŘđåŘŘïijŇ
- ä¡Ęåćđåď ğåŘŇæ■ěéŮť éŽŤåŘŕèČ¡äijŽéĂäæĹŘæIJĂçżĹèő■çżĈçš¿åžęçŽĎæ■§åď śãĂĆ [1]

äžěäÿŃäÿď äÿłç■ŰçŢěäżŐäÿ■åŘŇèğŠåžęèŕŢåŻ¿è¿¿åĹřæŻť åě¡çŽĎåźşèąąïijŽ

- post Local SGD åřĘèő■çżĈè£ĞçÍŃåĹĘæĹŘäÿď äÿłéŸűæőţïijŽçňňäÿĂéŸűæőţ trainers éŮť åŘŇæ■ěçŽĎéŮť éŽŤäÿž 1 äÿłæ■ěéŢ£ïijŇå■şåŘŇæ■ěSGDïijŇæİěä£ÌèŕĄæIJĂçżĹèő■çżĈçš¿åžęïijŽ HïijŇæİěæŘŘå■Ğèő■çżĈåŘđåŘŘãĂĆ
- Adaptive Communication Local SGD éĂŽè£ĞåŁÍæĂĄçŽĎèřĈæŢť åŔĆæŢřåŘŇæ■ěçŽĎéŮť éŽŤæİěåŕ

Fleet äÿ■åőđçŐřäžĘ `post Local SGD` åŠŇ `Adaptive Communication Local SGD` äÿď çğ■ç■ŰçŢěãĂĆ

åŁ§èČ¡æŢĹædIJ

åőđéłŇèő¿ç¡ő

model	dataset	local batch size	cluster	dtype	warming up	learning rate decay
resnet50	Imagenet	128	4 x 8 x V100	FP32	30	polynomial

åőđéłŇçżŞædIJ

local step	qps	acc1	acc5
1	8270.91	0.7579	0.9266
2	8715.67	0.7533	0.9265
4	8762.66	0.7551	0.9260
8	9184.62	0.7511	0.9239
16	9431.46	0.7429	0.9206
ADACOMM	8945.74	0.7555	0.9270

åŕŕäžěçIJŃåĹřåIJĺ post Local SGD ïijĹåŻžåőŽåŘŇæ■ěéŮťéŽŤïijĽæČĚåĘţäÿŃïijŇæŻťæŰřéŮťéŽŤèűł
å¡Şä¡£çŤĺ ADAPTIVE COMMUNICATION ç■ŰçŢěåŘŐïijŇèő■çżČåIJĺåŘđåŘŘåŠŇçš¿åžęéŮť è¿¿åĹřäžĘäÿ

ä¡£çŤĺæŰźæşŢ

äÿŃæŰĞåřĘäžěå■ŢæIJž8å■ąèő■çżČ ResNet50 äÿžä¿Ńå■ŘïijŇçőÅå■ŢäżŃçż■Local SGD çŽDçŤĺæşŢãÁĆéIJÄèęĄæşĺæĎŔçŽDæŸŕ å■ŢæIJžåĚńå■ąçŽDéÄŽä£ąéČ¡åIJĺåŘŇäÿĂæIJžåŹĺèŁĆçĆźå äÿĂèĹňæČĚåĘţäÿŃåŔĆæŢřåŘŇæ■ëäÿ■äijŽæĹŘäÿžèő■çżČçŽDçŞűéćĽïijŇè£ŹéĞŇåŔłæŸŕäžěåĚűäÿžä¿Ńå äÿ■ Local SGD åŔĆæŢřçŽDèő¿ç¡õãÁĆ

åőŽäźĽ'Local SGD çŻÿåĚşç■ŰçŢě

çŤĺæĹůéęŰåĚĹéIJÄèęĄåőŽäźĽ'paddle SGD åŕźèśąïijŇåźűåIJĺS-GDåŕźèśąäÿ■èő¿ç¡õå■ęäźäçŐĞåŔĆæŢřãÁĆçŻőåĽ'■local SGDåŠŇèĞłéĂĆåžŤæ■ëéŢ£ local SGDéČ¡äżĚæŤŕæŇĄSGDåŠŇMomentumäÿd'çğ■äijŸåŇŰåŹĺãÁĆ

- åIJĺ**post Local SGD** äÿ■ïijŇæIJĽäÿd'äÿłåŔĆæŢř `begin_step` åŠŇ `k_steps`ïijŇåśĂéČĺæŻť'æŰřåŠŇåŔĆæŢřåŘŇæ■ëéČ¡çŤśæąĘæđűèĞłåŁĺåőŇæĹŘãÁĆbegin_step æŇĞåőŽäżŐçňňåĞääÿłstepäźŃåŘŐè£ŻèąŇlocal SGDçőŮæşŢïijŇåŔŰåÄijäÿžåd'ğäžŐ0çŽDæŢť'æŢřïij æŇĞåőŽèő■çżČè£ĞçĺŃäÿ■çŽDåĚĺåśÄåŔĆæŢřæŻť'æŰřéŮť éŽŤïijŇåŔŰåÄijäÿžåd'ğäžŐ0çŽDæŢť'æŢ

```
strategy = fleet.DistributedStrategy()
strategy.localsgd = True
strategy.localsgd_configs = {
    "k_steps": 1,
    "begin_step": 1,
}
```

- åIJĺ **èĞłéĂĆåžŤæ■ëéŢ£ local SGD** äÿ■ïijŇæIJĽäÿd'äÿłåŔĆæŢř `begin_step` åŠŇ `init_k_steps`ãÁĆbegin_step æŇĞåőŽäżŐçňňåĞääÿłstepäźŃåŘŐè£ŻèąŇèĞłéĂĆåžŤlocal SGDçőŮæşŢïijŇåŔŰåÄijäÿžåd'ğäžŐ0çŽDæŢť'æŢřïijŻçŤĺæĹůéIJÄèęĄèő¿ç¡õinit_k_stepsä¡IJäÿžçňňäÿ äźŃåŘŐçŽDåŘŇæ■ëéŮť éŽŤårĘçŤśäÿŁæŰĞäÿ■çŽDåĚňåijŔåŁĺæÄĄçąőåőŽïijŇåIJĺå■ęäźäçŐĞè¿Čåd åd'Žè£ŻèąŇéĂŽä£ąäżŐèĂŇä£ÌèŕĄå£ńéĂ§æŤűæŢŻïijŻåIJĺå■ęäźäçŐĞè¿ČårŔæŮűïijŇåŔĆæŢřåŔŸåŇ åćđåd'ğstepïijŇåĞŔårŠéĂŽä£ąæňąæŢřïijŇäżŐèĂŇæŔŘå■Ğèő■çżČéĂ§åžęãÁĆ éIJÄèęĄæşĺæĎŔçŽDæŸŕåIJĺèĞłéĂĆåžŤæ■ëéŢ£ç■ŰçŢěäÿ■ïijŇçşżçż§äijŽéżŸèőd'éŹŘåĹűæIJÅåd'ğçŽ 16 stepïijŇå¡ŞåĚňåijŔèőąçőŮåĞžçŽDéŮť éŽŤåd'ğäžŐ16 æŰűïijŇæŇL'16 steps è£ŻèąŇåŔĆæŢřåŘŇæ■ëãÁĆ

```
strategy = fleet.DistributedStrategy()
strategy.adaptive_localsgd = True
strategy.adaptive_localsgd_configs = {
    "init_k_steps": 1,
    "begin_step": 1,
}
```

äÿŁè£řä¿Ńå■Řå■ŸæŤ¿åIJĺïijŽexample/resnet/train_fleet_static_localsgd.pyäÿŃéİćïijŇåĄĞèő¿èęĄè£ŘèąŇ2å■ąçŽĎäżżåŁąïijŇéĆčäźĹåŔłéIJĂåIJĺåŚ¡äżď'èąŇäÿ■æĽ'ğèąŇ:

```
fleetrun --gpus=0,1 train_fleet_static_overlap.py
```

æĆĺåřĘçIJŃåĹřæŸ¿çď'žåęĆäÿŃæŮěå£Ůä£ąæĄŕïijŽ

```
-----------  Configuration Arguments -----------
gpus: 0,1
heter_worker_num: None
heter_workers:
http_port: None
ips: 127.0.0.1
log_dir: log
...
------------------------------------------------
...
INFO 2021-01-18 22:01:11,969 launch_utils.py:472] Local start 2␣
↪processes. First process distributed environment info (Only For␣
↪Debug):
  ␣
↪+=======================================================================
   |                        Distributed Envs                       ␣
↪Value                      |
   +-----------------------------------------------------------------
↪----------------------+
   |                  PADDLE_CURRENT_ENDPOINT                  127.0.
↪0.1:10913               |
   |                      PADDLE_TRAINERS_NUM                       ␣
↪2                       |
   |                 PADDLE_TRAINER_ENDPOINTS         127.0.0.
↪1:10913,127.0.0.1:14758       |
   |                      FLAGS_selected_gpus                       ␣
↪0                       |
   |                        PADDLE_TRAINER_ID                       ␣
↪0                       |
  ␣
↪+=======================================================================
...
W0118 22:01:20.860090 45921 device_context.cc:362] Please NOTE:␣
↪device: 0, GPU Compute Capability: 7.0, Driver API Version: 10.2,␣
↪Runtime API Version: 9.2
W0118 22:01:20.864220 45921 device_context.cc:372] device: 0, cuDNN␣
↪Version: 7.4.
```

(äÿŃéąçżğçż■)

(çż■äÿŁéąţ)

```
W0118 22:01:25.578325 45921 gen_nccl_id_op_helper.cc:115] connect␣
↪addr=127.0.0.1:14758 failed 1 times with reason: Connection␣
↪refused retry after 0.5 seconds
[Epoch 0, batch 0] loss: 0.14602, acc1: 0.00000, acc5: 0.03125
[Epoch 0, batch 5] loss: 0.16445, acc1: 0.00000, acc5: 0.06250
```

6.4.5 èĞłåŁÍæůůåŘĹçš¿åžę

çőĂäżŃ

åIJÍä¡£çŤÍæŢřæ■őåźűèąŇåĹĘåÿČåijŕèő■çżČçŽĎåŘŇæŮű, æĹŚäžňè£ŸåŔŕäžěåijŢåĚěèĞłåŁÍæůůåŘĹçš¿åžę(Auto Mixed Precision) æİěè£ŻäÿĂæ■ěæŔŘå■ĞèŏçżČçŽĎéĂ§åžę.

äÿżæţĄçŽĎçěđçżŔç¡ŚçżIJæÍąåđŃéĂŽåÿÿä¡£çŤÍå■Ţçš¿åžę `single-precision (FP32)` æŢřæ■őæăijåijŔæİěå■ŸåĆÍæÍąåđŃåŔĆæŢřãĂĄè£ŻèąŇèő■çżČåŠŇéćĎæţŃ. åIJÍäÿŁè£řçŐŕèŁĆäÿ■ä¡£çŤÍå■Łçš¿åžę `half-precision (FP16)` æİěäžčæŻ£å■Ţçš¿åžę. åŔŕäžěåÿęæİěäžěäÿŃåě¡åd'Ď:

1. åĞŔåřŚåŕźGPU memory çŽĎéIJĂæśĆ: GPU æŸ¿å■Ÿäÿ■åŔŸæČĚåĘţäÿŃ, æŤŕæŇĄæŻt'åd'ğæÍąåđŃ / batch size
2. éŹ■ä¡ŐæŸ¿å■Ÿèŕžå£ŹæŮűçŽĎåÿęåő¡åŐŃåŁŻ
3. åŁăéĂ§GPU æŢřå■ęè£ŘçőŮéĂ§åžę (éIJĂèęĄGPU æŤŕæŇĄ[1])
4. GPUäÿŁ FP16 åŘđåŘŘæŸŕFP32 çŽĎ 2 - 8 åĂ■[2]

Paddle æŤŕæŇĄèĞłåŁÍæůůåŘĹçš¿åžęèőąçőŮ, åźűåőđçŐřäžĘ `èĞłåŁÍçżt'æŁd'FP32 ãĂĄFP16åŔĆæŢřåL'ŕæIJň`, `Dynamic loss scaling`, `opéżŚçŹ¡åŘ■å■Ţ` ç■Ľ'ç■ŰçŢěæİěéĄ£åĚ■ åŻă FP16 åŁÍæĂĄèŇČåŻt'è¿ČåřŔèĂŇåÿęæİěçŽĎæÍąåđŃæIJĂçżĹçš¿åžęæ■§åd'śãĂĆ Fleet ä¡IJäÿžPaddleéĂŽçŤÍçŽĎåĹĘåÿČåijŕèő■çżČAPIæŔŘä¿ŻäžĘçőĂå■ŢæŸŞçŤÍçŽĎæŐěåŔč, çŤÍæĹůåŔłéIJĂèęĄæůżåŁăăĞăèąŇäžčçăĄ åřśåŔŕåřĘèĞłåŁÍæůůåŘĹçš¿åžęåžŤçŤÍåĹřåŐ§æIJL'çŽĎåĹĘåÿČå

åŐ§çŘĘ

- TBA

æŞ■ä¡IJåőđèůţ

Fleet åřĘAMP åőđçŐřäÿž meta optimizer, çŤÍæĹůéIJĂèęĄæŇĞåőŽåĚűçŽĎ `inner-optimizer`. Fleet AMPæŤŕæŇĄæL'ĂæIJL' paddle optimziers åŠŇ FLeet meta otpimizers ä¡IJäÿžåĚű inner-optimizerãĂĆåŔłéIJĂèęĄåIJÍresetç¡ŚçżIJå§žçąĂäÿŁæL'ŞåijĂçŻÿåžŤçŽĎåijĂåĚşåŠŇéĚ■ç¡őçŻÿåžŤçŽĎéĂI

```
strategy = fleet.DistributedStrategy()
strategy.amp = True
strategy.amp_configs = {
    "init_loss_scaling": 32768,
    "decr_every_n_nan_or_inf": 2,
    "incr_every_n_steps": 1000,
    "incr_ratio": 2.0,
    "use_dynamic_loss_scaling": True,
    "decr_ratio": 0.5,
    "custom_white_list": [],
    "custom_black_list": [],
}
```

äÿŁè£řä¿Ńå■Řå■ŸæŤ¿åIJĺïijŽexample/resnet/train_fleet_static_amp.pyãĂĆ åĄĞèő¿èęĄè£ŘèąŇ8å■ąçŽĎäżżåŁąïijŇéĆčäźĹåŔłéIJĂåIJĺåŚ¡äžd'èąŇäÿ■æĽ'ğèąŇ:

```
fleetrun --gpus=0,1,2,3,4,5,6,7 train_fleet_static_amp.py
```

æĆĺåřĘçIJŃåĹřæŸ¿çd'žåęĆäÿŃæŮěå£Ůä£ąæĄŕïijŽ

```
-----------  Configuration Arguments -----------
gpus: None
heter_worker_num: None
heter_workers:
http_port: None
ips: 127.0.0.1
log_dir: log
...
------------------------------------------------
...
INFO 2021-01-19 14:46:03,186 launch_utils.py:472] Local start 8␣
↪processes. First process distributed environment info (Only For␣
↪Debug):
  ␣
↪+=======================================================================
   |                        Distributed Envs                              ␣
↪Value                      |
   +-----------------------------------------------------------------
↪----------------------+
   |                  PADDLE_CURRENT_ENDPOINT                   127.0.
↪0.1:54114                |
   |                      PADDLE_TRAINERS_NUM                          ␣
↪2                        |
   |                 PADDLE_TRAINER_ENDPOINTS  ... 0.1:24697,127.0.0.
↪1:53564,127.0.0.1:37181|
   |                      FLAGS_selected_gpus                          ␣
↪0                        |
   |                        PADDLE_TRAINER_ID                          ␣
↪0                        |
  ␣
↪+=======================================================================
```

(äÿŃéąçżğçż■)

(çż■äÿŁéąţ)

```
W0119 14:46:16.315114 84038 device_context.cc:362] Please NOTE:␣
↪device: 0, GPU Compute Capability: 7.0, Driver API Version: 10.2,␣
↪Runtime API Version: 9.2
W0119 14:46:16.320163 84038 device_context.cc:372] device: 0, cuDNN␣
↪Version: 7.4.
W0119 14:46:25.249166 84038 fuse_all_reduce_op_pass.cc:79] Find all_
↪reduce operators: 161. To make the speed faster, some all_reduce␣
↪ops are fused during training, after fusion, the number of all_
↪reduce ops is 8.
[Epoch 0, batch 0] loss: 0.19354, acc1: 0.00000, acc5: 0.00000
[Epoch 0, batch 5] loss: 0.20044, acc1: 0.00000, acc5: 0.00000
```

6.4.6 NV Dali Reader

- TBA

6.4.7 åĚűäżŰïijĹèřČèŁĆèţĎæžŘçŽĎéĚ■æŕŤãĂĄåćđåď'ğbsç■Ľ'ïijĽ'

åŐ§çŘĘ

PaddlePaddle ä¡£çŤĺâĂIJçž£çĺŃæśăâĂİæĺąåđŃèřČåžęåźűæĽ'ğèąŇOpïijŇOpåIJĺåŘŕåŁĺGPUèőąçőŮäźŃå
éĂŽåÿýéIJĂèęĄCPUçŽĎå■ŔåŁl'ïijŇçĎűèĂŇåęĆæđIJOpæIJňèžńå■ăçŤĺæŮűéÚt'å¿ĹårŔïijŇâĂIJçž£çĺŃæśăâ
æăźæ■őäžěå¿ĂçŽĎçżŔéłŇïijŇårźäžŐCPUäżżåŁąïijŇnum_threads=2 * dev_count
æŮűæĂğèČ¡è¿Čåě¡ïijŇ årźäžŐGPUäżżåŁąïijŇnum_threads=4 * dev_count
æŮűæĂğèČ¡è¿Čåě¡ãĂĆæşĺæĎŔïijŽçž£çĺŃæśääÿ■æŸŕèűŁåď'ğèűŁåě¡ãĂĆ

æŞ■ä¡lJåőđèůţ

çŤĺæĹůåŔłéIJĂèęĄæŇĞåőŽçŻÿåžŤçŽĎDistributedStrategy()çŽĎåijĂåĚşïijŇårśåŔŕäžěèő¿ç¡őçž£çĺŃæŢ

```
strategy = fleet.DistributedStrategy()

exe_strategy = paddle.static.ExecutionStrategy()
exe_strategy.num_threads = 3
strategy.execution_strategy = exe_strategy
```

äÿŁè£řä¿Ńå■Řå■ŸæŤ¿åIJĺïijŽexample/resnet/train_fleet_static_others.pyãĂĆ
åĄĞèő¿èęĄè£ŘèąŇ8å■ąçŽĎäżżåŁąïijŇéĆčäźĹåŔłéIJĂåIJĺåŚ¡äžď'èąŇäÿ■æĽ'ğèąŇ:

```
fleetrun --gpus=0,1,2,3,4,5,6,7 train_fleet_static_others.py
```

æĆĺårĘçIJŃåĹřæŸ¿çď'žåęĆäÿŃæŮěå£Ůä£ąæĄŕïijŽ

```
----------  Configuration Arguments ----------
gpus: None
heter_worker_num: None
heter_workers:
http_port: None
ips: 127.0.0.1
log_dir: log
...
-------------------------------------------------
...
INFO 2021-01-19 14:50:52,903 launch_utils.py:472] Local start 8␣
↪processes. First process distributed environment info (Only For␣
↪Debug):
   ␣
↪+=======================================================================
    |                        Distributed Envs                       ␣
↪Value                      |
    +-----------------------------------------------------------------
↪-----------------------+
    |                  PADDLE_CURRENT_ENDPOINT                  127.0.
↪0.1:20485               |
    |                      PADDLE_TRAINERS_NUM                      ␣
↪ 8                      |
    |                 PADDLE_TRAINER_ENDPOINTS  ... 0.1:23281,127.0.
↪0.1:41983,127.0.0.1:17503|
    |                      FLAGS_selected_gpus                      ␣
↪ 0                      |
    |                        PADDLE_TRAINER_ID                      ␣
↪ 0                      |
   ␣
↪+=======================================================================
...
W0119 14:51:04.500844 77798 device_context.cc:362] Please NOTE:␣
↪device: 0, GPU Compute Capability: 7.0, Driver API Version: 10.2,␣
↪Runtime API Version: 9.2
W0119 14:51:04.506238 77798 device_context.cc:372] device: 0, cuDNN␣
↪Version: 7.4.
W0119 14:51:12.378418 77798 fuse_all_reduce_op_pass.cc:79] Find all_
↪reduce operators: 161. To make the speed faster, some all_reduce␣
↪ops are fused during training, after fusion, the number of all_
↪reduce ops is 5.
[Epoch 0, batch 0] loss: 0.11252, acc1: 0.03125, acc5: 0.06250
[Epoch 0, batch 5] loss: 0.11252, acc1: 0.03125, acc5: 0.06250
[Epoch 0, batch 10] loss: 0.11252, acc1: 0.03125, acc5: 0.06250
[Epoch 0, batch 15] loss: 0.11252, acc1: 0.03125, acc5: 0.06250
```

6.5 åď'ğæÍąåđŃèő■çżČäijŸåŇŰ

6.5.1 Forward Recomputation Backpropagation

çőĂäżŃ

ä¡£çŤÍæŻť'åď'ğçŽĎæÍąåđŃåŠŇæŻť'åď'ğçŽĎbatch size åŕŕäžěåÿęæİěæŻť'åě¡çŽĎæŢĹæđIJïijŇä¡EçŤÍæĹ
Fleet æŔŘä¿Żäžěäÿŋäÿď'çğ■ç■ŰçŢěæİèèğčåĘşåď'ğæÍąåđŃïijÍåď'ğbatch
sizeïijL'èő■çżČäÿ■åŕŕèČ¡éĄĞåĹřçŽĎæŸ¿å■ŸçŞűéćĹéŮőécŸïijŽ

Forward Recomputation BackpropagationïijĹFRBïijL'

èŕěç■ŰçŢěéĂŽè£ĞæÿĚéŹď'æ■čåŘŚèőąçőŮè£ĞçÍŃäÿ■çŽĎäÿ■éŮť'èőąçőŮçżŞæđIJïijŇæİěéŹ■ä¡Őèő■çż
Size çŽĎèő■çżČãĂĆ

Recompute-Offload

å§žäžŐRecompute ç■ŰçŢěïijŇåřĘæŸ¿å■Ÿäÿ■checkpoint å■ÿè¡¡åĹřHost
åĘĚå■Ÿäÿ■ïijŇè£ŻäÿĂæ■ěèŁĆçIJĄæŸ¿å■Ÿçl'žéŮť'æŤŕæŇĄæŻť'åď'ğbatch
SizeçŽĎèő■çżČãĂĆ

åŐ§çŘĘ

Recomputation

æĹŚäžñç§ěéĄŞïijŇæůśåžęå■ęäźăç¡ŚçżIJçŽĎäÿĂæňąèő■çżČè£■äžčåŇĚåŘńäÿL'äÿłæ■ěéłď'ïijŽ

- **åL'■åŘŚèőąçőŮïijŽ** è£ŘèąŇåL'■åŘŚçőŮå■Ř(Operator) æİěèőąçőŮäÿ■éŮť'éŽŘåśĆ(Variable) çŽĎåĂij ãĂĆ
- **åŕ■åŘŚèőąçőŮïijŽ** è£ŘèąŇåŕ■åŘŚçőŮå■ŘæİěèőąçőŮåŔĆæŢř(Parameter)çŽĎæćŕåžęãĂĆ
- **äijŸåŇŰïijŽ** åžŤçŤÍäijŸåŇŰçőŮæşŢäžěæŻť'æŰřåŔĆæŢřåĂij ãĂĆ

åIJÍåL'■åŘŚèőąçőŮè£ĞçÍŃäÿ■ïijŇåL'■åŘŚçőŮå■ŘäijŽèőąçőŮåĞžåď'ğéĞŔçŽĎäÿ■éŮť'çżŞæđIJïijŇçŤšä
bize èűŁåď'ğïijŇäÿ■éŮť'çżŞæđIJå■ăçŤÍçŽĎåĘĚå■Ÿäź§åřśèűŁåď'ğãĂĆéčđæąÍæäÿå£ČæąĘæđűäijŽä¡£çŤÍ
Variableæİěå■ŸåĆÍè£ŹäžŻéŽŘåśĆçŽĎäÿ■éŮť'çżŞæđIJãĂĆå¡ŞæÍąåđŃåśĆæŢřåŁăæůśæŮűïijŇåĚűäÿ■éŮť'çż
å■ăæ■őåď'ğéĞŔçŽĎåĘĚå■ŸãĂĆéčđæąÍæäÿå£ČæąĘæđűçŽĎæŸ¿å■ŸåŽđæŤűæIJžåĹűäijŽåŘŁæŮűæÿĚéŹď'
ä¡ĘæŸŕæIJL'äžŻäÿ■éŮť'çżŞæđIJæŸŕåŔ■åŘŚèőąçőŮè£ĞçÍŃäÿ■çőŮå■ŘçŽĎè¿ŞåĚěïijŇè£ŹäžŻäÿ■éŮť'çżŞæ

åŕźäžŐåď'ğåřŘåŽžåőŽçŽĎåĘĚå■Ÿæİèèŕť'ïijŇåęĆæđIJçŤÍæĹůåÿŇæIJŻä¡£çŤÍåď'ğbatch
bize çŽĎæŢřæ■őè£ŻèąŇèő■çżČïijŇåĹŹåřĘåřijèĞť'å■Ţäÿłäÿ■éŮť'çżŞæđIJå■ăçŤÍåĘĚå■Ÿåćđåď'ğïijŇéĆčäźĹ

FRBæŸŕåřĘæůśåžęå■ęäźăç¡ŚçżIJåĹĞåĹĘäÿžkäÿłéČÍåĹĘïijÍsegmentsïijL'ãĂĆåŕźæŕŘäÿłsegment
èĂŇèÍĂïijŽåL'■åŘŚèőąçőŮæŮűïijŇéŹď'äžĘåřŘéČÍåĹĘå£Ěéążå■ŸåĆÍåIJÍåĘĚå■Ÿäÿ■çŽĎVariable
åď'ŰïijŇåĚűäżŰäÿ■éŮť'çżŞæđIJéČ¡åřĘèćńåĹăéŹď'ïijŽåIJÍåŔ■åŘŚèőąçőŮäÿ■ïijŇéęŰåĚĹéĞ■æŰřèőąçőŮäÿ

åŠŇæŹőéĂŽçŽĎç¡ŚçżIJè£■äžčçŻÿæŕŤïijŇåď'Žèőąçő̊ŮäžĘäÿĂéĄ■åL'■åŘŚçő̊Ůå■ŘãĂĆ
åĚůä¡Şè£ĞçÍŃåęĆäÿŃåŻ¿æL'Ăçď'žïijŽ

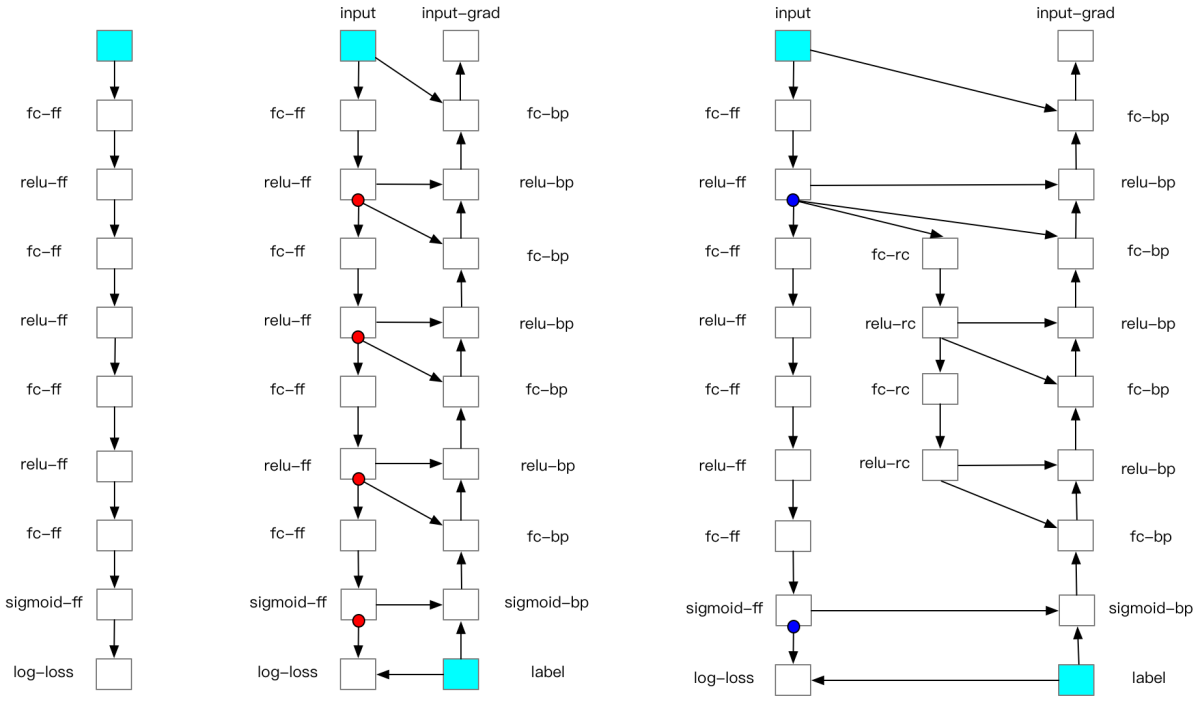

æĹŚäžňæŁŁåĹĞåĹĘç¡ŚçżIJçŽĎåŔŸéĞŔåŔńåĄŽcheckpointsãĂĆ
éĆčäźĹéŮőécŸæİěäžĘïijŇåęĆä¡ŢéĂL'æŃl'checkpoints åŚćïij§èĞłäžŐFRBæŰźæşŢæŔŘåĞžäžěæİěïijŇåď'gé
æĹŚäžňç§ěéĄŞæůšåžęå■ęäźăç¡ŚçżIJéĂŽåÿÿæŸŕçŤśäÿĂäÿłäÿłæĺąåİŮäÿšèĄŤå¿ŮåĹrçŽĎïijŇærŤåęĆResNet-
50çŤś16äÿłblockäÿšèĄŤèĂŇæĹŘïijŇ Bert-LargeçŤś24äÿłEncoder layers
äÿšèĄŤèĂŇæĹŘïijŇäžěäÿd'äÿłå■ŘæĺąåİŮäÿ■éŮťçŽĎåŔŸéĞŔä¡IJäÿžåĹĞåĹĘçĆźåŕśæŸŕäÿĂäÿłå¿Ĺåě¡çŽĎ
åŕźäžŐéİđäÿšèĄŤçŽĎç¡ŚçżIJïijĹæŕŤåęĆåŘńæIJL'åď'géĞŔshortcutçżŞædĎçŽĎç¡ŚçżIJïijL'ïijŇFRBäź§æŤŕæ
åŘłæŸŕåŔŕèČ¡åď'ŽèĂŮèť'źäÿĂçĆźåĘĚå■ŸïijĹçŤĺäžŐå■ŸåĆĺshortcutçŽĎVariableïijL'ãĂĆ

Recompute-Offload

åIJĺäÿŁéİćçŽĎRecomputation æ■ěéłd'äÿ■ïijŇåŘŇæăůä¡IJäÿžForward
äÿ■éŮť'çżŞædIJçŽĎcheckpoints äijŽél'żçŢŹæŸ¿å■ŸïijŇæŰźä¿£åIJĺBackward
äÿ■éĞ■èőąçő̊ŮãĂĆ çĎűèĂŇåIJĺcheckpoint çŽĎçŤ§åŚ¡åŚĺæIJ§äÿ■ïijŇäż■æIJL'äÿĂæőţè¿ČéŢ£çŽĎæIJłèćńä¡
è£Źäź§æŸŕåŕźæŸ¿å■ŸçŽĎäÿĂçğ■æţłèť'źãĂĆ Recompute-Offload
åŐ§çŘĘåď'gèĞť'åŔŕäžěåĹĘäÿžäÿd'æ■ěïijŽ

- ForwardïijŽ å¡Şcheckpointå IJĺåL'■åŘŚäÿ■èćńçŤ§æĹŘåŘŐïijŇářĘåĚűå■ÿè¡¡(Offload)åĹřHost åĘĚå■Ÿäÿ■ïijŇèől'åĚűæL'Ăå■ăæ■őçŽĎæŸ¿å■ŸåŔŕäžěèćńéĞŁæŤ¿ãĂĆ
- BackwardïijŽå¡ŞcheckpointåIJĺåŔ■åŘŚäÿ■èćńéĞ■æŰřèřČçŤĺäźŃåL'■ïijŇářĘåĚűécĎåŔŰ(Pre-fetch) åŽđæŸ¿å■Ÿäÿ■ïijŇåőŇæĹŘäźŃåŘŐçŽĎéĞ■èőąçő̊ŮãĂĆ

æşĺæĎŔïijŽ

- åŽääÿžcheckpoint åIJĺåĘĚå■ŸåŠŇæŸ¿å■ŸéŮť'çŽĎæŃůèť'Ìè¿ČæĚćïijŇèŕěç■ŰçŢěæŸŕéĂŽè£Ğè£ŽäÿĂ
sizeïijŇ éIJĂèęĄçŤĺæĹůæİĆèąąèő■çżČåŘđåŘŘåŠŇbatch size ãĂĆ
- Recompute-Offload æŤŕæŇĄåd'Žå■ąåźűèąŇèő■çżČïijŇ

å¡Şåď'Žå■ąåźűèąŇæŮűåijĂåŘŕOffloadïijŇèő■çżČäÿ■åŘŇäÿĂèŁĆçĆźäÿŁæL'ĂæIJL'GPU äÿŁçŽĎcheckpoints éČ¡ářĘå■ÿè¡¡åĹřHost åĘĚå■Ÿäÿ■ïijŇäijŽå■ŸåIJĺäžěäÿŇéčŐéŹl'ïijŽ

- PCIe åÿęåő¡çŞűéćĹïijŽ åŘŇäÿĂèŁĆçĆźäÿŁçŽĎæL'ĂæIJL'GPU åŠŇHost åĘĚå■ŸéŮt'åĚśäžńäÿĂæăźPCIe åÿęåő¡ïijŇåęĆåŘŇäÿĂèŁĆçĆźäÿŁGPU æŢřéĞŔè¿Čåď'ŽïijĹå■ŢæIJžåĚńå■ąïijL'åőźæŸŞåŻääÿžPCIe åÿęåő¡éŹŘåĹűèől'è■çżČéĂ§åžęè£ŻäÿĂæ■ěåĞŔæĚć
- Host åĘĚå■ŸæžćåğžïijŽ å¡ŞåŘŇäÿĂèŁĆçĆźäÿŁGPU æŢřéĞŔè¿Čåď'ŽïijŇäÿŤærŔåijăGPU checkpoints size è¿Čåď'ğæŮűïijŇéIJĂèęĄæşÍæĎŔå■ÿè¡¡éĞŔæŸŕåŘęèűĚåĞžHost åĘĚå■Ÿåď'ğářŔãĂĆ

æŢĹædIJ

æĹŚäžňåIJíBERT-LargeæÍąåđŇäÿŁåŕźRecompute çŽĎæŢĹædIJè£ŻèąŇäžĘæţŇèŕŢïijŇRecompute åŔŕäžěèől'batch size æL'l'åď'ğ 10åĂ■ïijŇ Offload åŔŕäžěåIJíRecompute çŽĎå§žçąĂäÿŁåĘ■æL'l'åď'ğ1.43 åĂ■ãĂĆ batch size = #seq * seq_max_len çąňäžű: å■Ţå■ą V100 32GB

ç■ŰçŢě	amp	amp + Recompute	amp + Recompute + offload
batch size	18 * 512	180 * 512	258 * 512
speed	23.94 sents/s	17.82 sents/s	15.47 sents/s

ä¡£çŤÍæŰźæşŢ

äÿžäžĘä¡£çŤÍRecomputeç■ŰçŢěïijŇæĹŚäžňåřĘdist_strategy.recomputeèő¿ç¡őäÿžTrue åźűèő¿ç¡őæĹŚäžňäžŇåĚĹåőŽäźL'åě¡çŽĎcheckpointsãĂĆ checkpoint çŽĎéĂL'åŔŮåŔŕäžěåŔĆèĂČèőžæŰĞ ãĂŁTraining Deep Nets with Sublinear Memory CostãĂŃ ãĂĆ

çd'žä¿Ńäÿ■ä¡£çŤÍçŽĎResNet50 æÍąåđŇçŽĎ checkpoint äÿ■æŸŕåŻžåőŽçŽĎïijŇäÿ■çňęåŘĹ Offload çŽĎèęĄæśĆïijŇåŻžèŕěåŁ§èČ¡æŽĆæŮăæşŢåijĂåŘŕãĂĆ å¡Şä¡£çŤÍ Transformer æŮűïijŇåŔŕäžěéĂL'åŔŮærŔäÿĂlayer çŽĎFC output ä¡IJäÿžcheckpointïijŇ è£ŹæŮűåŘĎäÿłlayer çŽĎcheckpoints shapes äÿĂèĞt'ïijŇåŔŕäžěä¡£çŤÍOffloadãĂĆ

res2a.add.output.5.tmp_0 ç■L'æŸŕçŤÍæĹůçżĎç¡ŚæŮűåőŽäźL'çŽĎ variable name

```
checkpoint_idx = ["2a", "2b", "2c", "3a", "3b", "3c", "3d", "4a",
↪"4b", "4c", "4d", "4e", "4f", "5a", "5b", "5c"]
checkpoints = ['res{}.add.output.5.tmp_0'.format(idx) for idx in␣
↪checkpoint_idx]
strategy = fleet.DistributedStrategy()
strategy.recompute = True
strategy.amp = True
strategy.recompute_configs = {
    "checkpoints": checkpoints,
```

(äÿŃéąçżğçż■)

(çż■äÿŁéąţ)

```
    "enable_offload": False,
    "checkpoint_shape": []
    }
```

äÿŁè£řä¿Ńå■ŘçŽĎåőŇæŢťäžčçăĄå■ŸæŤ¿åIJĺïijŽtrain_fleet_recompute.pyäÿŃéİćãĂĆåĄĞèő¿èęĄè£Řè

```
fleetrun --gpus=0,1 train_fleet_recompute.py
```

æĆĺåřĘçIJŃåĹřæŸ¿çd'žåęĆäÿŇæŮěå£Ůä£ąæĄŕïijŽ

```
----------  Configuration Arguments ----------
gpus: 0,1
heter_worker_num: None
heter_workers:
http_port: None
ips: 127.0.0.1
log_dir: log
...
-----------------------------------------------
...
   ⌊
 ↪+=========================================================================
    |                           Distributed Envs                          ⌊
 ↪Value                       |
    +-------------------------------------------------------------------
 ↪-----------------------+
    |                   PADDLE_CURRENT_ENDPOINT                   127.0.
 ↪0.1:17901              |
    |                       PADDLE_TRAINERS_NUM                        ⌊
 ↪ 2                      |
    |                  PADDLE_TRAINER_ENDPOINTS         127.0.0.
 ↪1:17901,127.0.0.1:18846       |
    |                       FLAGS_selected_gpus                        ⌊
 ↪ 0                      |
    |                         PADDLE_TRAINER_ID                        ⌊
 ↪ 0                      |
   ⌊
 ↪+=========================================================================
...
   ⌊
 ↪+=========================================================================
    |                                                                   ⌊
 ↪                |
    |                           DistributedStrategy Overview            ⌊
 ↪                |
    |                                                                   ⌊
 ↪                |
   ⌊
 ↪+=========================================================================
```

(äÿŃéąţçżğçż■)

(çż■äÿŁéąţ)

```
    |                              amp=True <-> amp_configs                 ␣
↪                |
    +-------------------------------------------------------------------
↪---------------+
    |                        init_loss_scaling                    32768.0␣
↪                |
    |                      incr_every_n_steps                       1000 ␣
↪                |
    |                 decr_every_n_nan_or_inf                        2    ␣
↪                |
    |                              incr_ratio                       2.0  ␣
↪                |
    |                              decr_ratio              0.
↪800000011920929            |
    |                 use_dynamic_loss_scaling                     True ␣
↪                |
  ␣
↪+==================================================================================
    |                        recompute=True <-> recompute_configs       ␣
↪                |
    +-------------------------------------------------------------------
↪---------------+
    |                              checkpoints           res2a.add.
↪output.5.tmp_0        |
    |                                                    res2b.add.
↪output.5.tmp_0        |
    |                                                    res2c.add.
↪output.5.tmp_0        |
    |                                                    res3a.add.
↪output.5.tmp_0        |
    |                                                    res3b.add.
↪output.5.tmp_0        |
    |                                                    res3c.add.
↪output.5.tmp_0        |
    |                                                    res3d.add.
↪output.5.tmp_0        |
    |                                                    res4a.add.
↪output.5.tmp_0        |
    |                                                    res4b.add.
↪output.5.tmp_0        |
    |                                                    res4c.add.
↪output.5.tmp_0        |
    |                                                    res4d.add.
↪output.5.tmp_0        |
    |                                                    res4e.add.
↪output.5.tmp_0        |
    |                                                    res4f.add.
↪output.5.tmp_0        |
    |                                                    res5a.add.
↪output.5.tmp_0        |
```

(äÿŃéąţçżğçż■)

(çż■äÿŁéąţ)

```
    |                                                    res5b.add.
↪output.5.tmp_0       |
    |                                                    res5c.add.
↪output.5.tmp_0       |
    |                          enable_offload                    False ␣
↪                |
   ␣
↪+=========================================================================
...
W0104 17:59:19.018365 43338 device_context.cc:342] Please NOTE:␣
↪device: 0, GPU Compute Capability: 7.0, Driver API Version: 10.2,␣
↪Runtime API Version: 9.2
W0104 17:59:19.022523 43338 device_context.cc:352] device: 0, cuDNN␣
↪Version: 7.4.
W0104 17:59:23.193490 43338 fuse_all_reduce_op_pass.cc:78] Find all_
↪reduce operators: 161. To make the speed faster, some all_reduce␣
↪ops are fused during training, after fusion, the number of all_
↪reduce ops is 5.
[Epoch 0, batch 0] loss: 0.12432, acc1: 0.00000, acc5: 0.06250
[Epoch 0, batch 5] loss: 1.01921, acc1: 0.00000, acc5: 0.00000
...
```

åőŇæŢť2å■ąçŽĎæŮěå£Ůä£ąæĄŕäź§åŔŕåIJĺ./log/çŽőå¡ŢäÿŃæ§ěçIJŃãĂĆäžĘèğčæŻťåď`Žfleetru
åŘŕåŁĺåĹĘåÿČåijŔäżżåŁąãĂĆ

6.5.2 Gradient Merge

çőĂäżŃ

äÿžäžĘæŔŘå■ĞæĺąåđŃçŽĎæĂğèČ¡ïijŇäžžäžňåijĂåğŃè£¡æśĆïijŽæŻť'åď'gèğĎæĺąçŽĎæŢřæ■őéŻĘãĂA
èŁŕçL'ĞåžŤè£ŘèĂŇçŢ§ãĂĆä¡ĘåIJĺåĹĘåÿČåijŔèő■çżČäÿ■ïijŇçżŔåÿÿäijŽéĄĞåĹřæŸ¿å■ŸæĹŰèĂĚåĘĚå■Ÿ

- è¿ŞåĚěçŽĎæŢřæ■õè£Ğåď'gïijŇä¿ŃåęĆèğĘéćŚçśżèő■çżČæŢřæ■őãĂĆ
- æůśåžęæĺąåđŃçŽĎåŔĆæŢřè£Ğåď'ŽæĹŰè£Ğåď'gïijŇæL'ĂéIJĂçŽĎå■ŸåĆĺçl'žéŮt'èűĚåĞžäžĘåĘĚå■Ÿ
- AIèŁŕçL'ĞçŽĎåĘĚå■ŸæIJL'éŹŘãĂĆ

äÿžäžĘèČ¡æ■čåÿÿåőŇæĹŘèő■çżČïijŇæĹŚäžňéĂŽåÿÿåŔłèČ¡ä¡£çŤĺè¿ČåřŔçŽĎbatch
size äžěéŹ■ä¡ŐæĺąåđŃèő■çżČäÿ■çŽĎæL'ĂéIJĂèęĄçŽĎå■ŸåĆĺçl'žéŮt'ïijŇè£ŹåřĘåŕijèĞt'å¿Ĺåď'ŽæĺąåđŃæŮ
size æĬěæŔŘéńŸæĺąåđŃçŽĎçš¿åžęãĂĆ

Gradient Merge ç■ŰçŢěçŽĎäÿżèęĄæĂİæČşæŸŕåřĘè£đçż■åď'Žäÿłbatch
æŢřæ■őèő■çżČå¿ŮåĹřçŽĎåŔĆæŢřæćŕåžęåŘĹåźűåĄŽäÿĂæňąæŻť'æŰřãĂĆ
åIJĺèŕěèő■çżČç■ŰçŢěäÿŃïijŇèŹ¡çĎűäżŐå¡ćåijŔäÿŁçIJŃä¿İçĎűæŸŕåřŔbatch
èğĎæĺąçŽĎæŢřæ■őåIJĺèő■çżČïijŇä¡ĘæŸŕæŢĹæđIJäÿŁåŔŕäžěè¿¿åĹřåď'Žäÿłåřŕbatch
æŢřæ■őåŘĹåźűæĹŘåď'ğbatch åŘŐèő■çżČçŽĎæŢĹæđIJãĂĆ

åŐ§çŘĘ

Gradient Merge åŔłæŸŕåIJÍèő■çżČæţĄçÍŃäÿŁåĄŽäžĘäÿĂäžŻå¿őěřČïijŇè¿¿åĹřæĺąæŃ§åĞžåď ğbatch size èő■çżČæŢĹæđIJçŽĎçŽőçŽĎãĂĆåĚůä¡Şæİěèŕť ïijŇåřśæŸŕä¡£çŢĺèŃěáźšåŐ§æIJL'åď ğåřŔçŽĎbatch æŢřæ■őè£ŻèąŇèő■çżČïijŇå■şéĂŽè£ĞâĂIJåL'■åŘŚ+åŔ■åŘŚâĂİ ç¡ŚçżIJèőąçőŮå¿ŮåĹřæćŕåžęãĂĆåĚűéŮť äijŽæIJL'äÿĂéČĺåĹĘæŸ¿å■Ÿ/åĘĚå■ŸçŢĺäžŐå■ŸæŤ¿æćŕåžęïijŇçİ æŢřæ■őèő■çżČçŽĎæŢĹæđIJãĂĆ

åIJÍè¿Čåď ğçŽĎçšŠåžęäÿŁçIJŃïijŇ GM æŸŕåřĘèő■çżČäÿĂäÿłstep çŽĎè£ĞçÍŃçŢśåŐ§æİěçŽĎ âĂIJåL'■åŘŚ + åŔ■åŘŚ + æŻť'æŰřâĂİ æŤźåŘŸæĹŘ âĂIJïijĹåL'■åŘŚ + åŔ■åŘŚ + æćŕåžęçţ'ŕåŁäïijL'x k + æŻť'æŰřâĂİïijŇ éĂŽè£ĞåIJÍæIJĂçžĹæŻť'æŰřåL'■è£ŻèąŇ k æňąæćŕåžęçŽĎçţ'ŕåŁäæĺąæŃ§åĞž batch size æL'l'åď ğ k åĂ■çŽĎæŢĹæđIJãĂĆ æŻť'åĚůä¡ŞçżĘèŁĆåŔŕäžěåŘĆèĂČ ãĂŁMG-WFBP: Efficient Data Communication for Distributed Synchronous SGD AlgorithmsãĂŃ ãĂĆ

ä¡£çŤÍæŰźæşŢ

Gradient Merge ç■ŰçŢěåIJÍä¡£çŤÍæŰźéÍćäź§å¿ĹçőĂå■ŢïijŇçŤÍæĹůåŘłéIJĂèęĄåőŽäźL'ărĘåď ŽářŚbatch çŽĎæŢřæ■őèőąçőŮåĞžçŽĎæćŕåžęåŘăåŁäæŻť'æŰřæĺąåđŃåŘĆæŢřïijŇä¿£åŔŕäžěåőđçŐřåď ğbatch èő■çżČçŽĎçŽőçŽĎãĂĆ

èő■çżČäžčçăĄçŽĎæąĘæđűåŠŃåĚűäžŰfleet èő■çżČäžčçăĄå§žæIJňäÿĂæăůïijŇçŤÍæĹůåŘłéIJĂèęĄåIJÍ fleet.DistributedStrategy äÿ■éĚ■ç¡őGradient Merge çŻÿåĚşåŘĆæŢřå■şåŘŕãĂĆ

åĄĞèő¿æĹŚäžňåőŽäźL'äžĘbatch size äÿž NïijŻéĂŽè£Ğèő¿ç¡ők_stepsïijŇä¡£çŤÍ4äÿłbatch sizeæİěæĺąæŃ§äÿĂäÿłåď ğbatchçŽĎèő■çżČïijŇäžŐèĂŃè¿¿åĹřäžĘbatch size äÿž 4*N çŽĎèő■çżČæŢĹæđIJãĂĆ

åIJÍgradient_merge_configsäÿ■ïijŇavg éĂL'éąźçŤÍäžŐæŐğåĹűæćŕåžęçţ'ŕèőąçŽĎå¡ćåijŔïijŽå¡Ş True æŰűïijŇäijŽåŕźæŕŔæňąçŽĎæćŕåžęæśĆåŠŃåźűåĄŽåźşåİĞïijŻåŔ■ázŃåřĘçŻť'æŐěåŕźæćŕåžęæśĆåŠŃïijŇ

```
strategy = fleet.DistributedStrategy()
# ä¡£çŤÍGradient mergeç■ŰçŢěåźűèő¿ç¡őçŻÿåĚşåŘĆæŢř
strategy.gradient_merge = True
strategy.gradient_merge_configs = {"k_steps": 4, "avg": True}
```

äÿŁè£ŕä¿Ńå■ŘçŽĎåőŇæŢť'äžčçăĄå■ŸæŤ¿åIJÍïijŽtrain_fleet_gradient_merge.pyäÿŃéÍćãĂĆåĄĞèő¿èęA

```
fleetrun --gpus=0,1 train_fleet_gradient_merge.py
```

æĆĺåřĘçIJŃåĹřæŸ¿çď'žåęĆäÿŇæŮěå£Ůä£ąæĄŕïijŽ

```
----------  Configuration Arguments ----------
gpus: 0,1
heter_worker_num: None
heter_workers:
http_port: None
ips: 127.0.0.1
log_dir: log
...
```

(äÿŃéąţçżğçż■)

(çż■äÿŁéąţ)

```
-------------------------------------------------
...
   ␣
↪+=========================================================================
    |                          Distributed Envs                           ␣
↪Value                       |
    +-------------------------------------------------------------------
↪-----------------------+
    |                  PADDLE_CURRENT_ENDPOINT                  127.0.
↪0.1:17901               |
    |                      PADDLE_TRAINERS_NUM                          ␣
↪ 2                      |
    |                 PADDLE_TRAINER_ENDPOINTS         127.0.0.
↪1:17901,127.0.0.1:18846       |
    |                      FLAGS_selected_gpus                          ␣
↪ 0                      |
    |                        PADDLE_TRAINER_ID                          ␣
↪ 0                      |
   ␣
↪+=========================================================================
...
   ␣
↪+=========================================================================
    |                                                                     ␣
↪               |
    |                          DistributedStrategy Overview               ␣
↪               |
    |                                                                     ␣
↪               |
   ␣
↪+=========================================================================
    |                 gradient_merge=True <-> gradient_merge_configs      ␣
↪               |
    +-------------------------------------------------------------------
↪--------------+
    |                                 k_steps                      4      ␣
↪               |
    |                                     avg                     True   ␣
↪               |
   ␣
↪+=========================================================================
...
W0104 17:59:19.018365 43338 device_context.cc:342] Please NOTE:␣
↪device: 0, GPU Compute Capability: 7.0, Driver API Version: 10.2,␣
↪Runtime API Version: 9.2
W0104 17:59:19.022523 43338 device_context.cc:352] device: 0, cuDNN␣
↪Version: 7.4.
W0104 17:59:23.193490 43338 fuse_all_reduce_op_pass.cc:78] Find all_
↪reduce operators: 161. To make the speed faster, some all_reduce␣
↪ops are fused during training, after fusion, the number of all_
↪reduce ops is 5.
```

(äÿŇeąţężğçż■)

(çż■äÿŁéąţ)

```
[Epoch 0, batch 0] loss: 0.12432, acc1: 0.00000, acc5: 0.06250
[Epoch 0, batch 5] loss: 1.01921, acc1: 0.00000, acc5: 0.00000
...
```

åőŇæŢť'2å■ąçŽĎæŮěå£Ůä£ąæĄŕäź§åŔŕåIJĺ./log/çŻőå¡ŢäÿŃæ§ěçIJŃãĂĆäžĘèğčæŻť'åď'Žfleetru åŘŕåŁĺåĹĘåÿČåijŔäżżåŁąãĂĆ

6.5.3 ä¡£çŤĺLARS / LAMB äijŸåŇŰåĹĘåÿČåijŔèűĚåď'gbatch èő■çżČ

çőĂäžŃ

åIJĺæŢřæ■őåźűèąŇåĹĘåÿČåijŔèő■çżČåIJžæŹŕäÿ■, åÿÿä¡£çŤĺåćđåŁăGPUæŢřéĞŔçŽĎæŰźåijŔæİěåŁăéäÿžäžĘä£ÌèŕĄGPUçŽĎçőŮåŁŻå¿ŮåĹŕåĚĚåĹĘåĹl'çŤĺ, æŕŔåijăGPUå■ąäÿŁçŽĎ-batch sizeéČ¡éIJĂèęĄèűşåď'§åď'ğãĂĆ åŻăæ■ď'åIJĺåćđåŁăGPU æŢřéĞŔåŘŇæŮű, èő■çżČçŽĎåĚĺåśĂbatch size äź§äijŽåŔŸåď'ğãĂĆ

ä¡ĘèűŁåď'ğçŽĎåĚĺåśĂbatch size äijŽåÿęæİěèő■çżČçŽĎæŤűæŢŻéŮőécŸ[1] [2]:

- æĺąåđŃæIJĂçżĹçš¿åžęæ■§åď'ś
- æŤűæŢŻéĂ§åžęåŔŸæĚć, éIJĂèęĄæŻť'åď'ŽçŽĎepoch æĽ'■èČ¡æŤűæŢŻ

LARS[3] åŠŇ LAMB[4] äÿď'äÿłäijŸåŇŰç■ŰçŢěåÿÿçŤĺæİěèğčåĘşäÿŁè£řèűĚåď'gbatch èő■çżČäÿ■çŽĎæŤűæŢŻéŮőécŸãĂĆ

Paddle åőđçŐřäžĘè£Źäÿď'çğ■äijŸåŇŰç■ŰçŢěïijŇpaddle.distributed.fleet ä¡IJäÿžPaddleéĂŽçŤĺçŽĎåĹĘåÿČåijŔèő■çżČAPIæŔŘä¿ŻäžĘçőĂå■ŢæŸŞçŤĺçŽĎæŐěåŔč, çŤĺæĹůåŘłéIJĂèęĄæůżåŁäåĞäèąŇäžčçăĄåŕšåŔŕåřĘç■ŰçŢěåŁäåĚěåĹřåŐ§æIJL'çŽĎèő■çżČäÿ■ãĂĆ éĂŽè£Ğè£Źäÿď'äÿłäijŸåŇŰç■ŰçŢě, æĹŚäžňåIJĺèűĚåď'gbatch åIJžæŹŕäÿ■åőđçŐřäžĘæŻť'å£ńçŽĎæŤűæŢŻéĂ§åžęåŠŇæŮäæ■§çŽĎçš¿åžę, çżŞåŘĹFleet äÿ■åĚűäżŰçŽĎç■ŰçŢě(e.g. AMP) åŔŕäžěçijl'ç§■æŢť'ä¡Şèő■çżČæŤűæŢŻæŮűéŮť'ãĂĆ

åŐ§çŘĘ

LARS

LARS åĚňåijŔåęĆäÿŃïijŽ

$$local\_learning\_rate = learning\_rate * lars\_coeff * \frac{||param||}{||gradient|| + lars\_weight\_decay * ||param||}$$

$$velocity = mu * velocity + local\_learning\_rate * (gradient + lars\_weight\_decay * param + epsilon)$$

$$param = param - velocity$$

åŔŕäžěçIJŃåĹřLARS åĚűåőđæŸŕåIJĺ åÿęweight decay çŽĎmomentum äijŸåŇŰåŹĺçŽĎå§žçąĂäÿŁåŁäåĚěäžĘlocal learning rate çŽĎéĂżè¿Ś, åŕźæŕŔäÿĂåśĆçŽĎlearning rate è£ŻèąŇäžĘæŤ¿çijl'ãĂĆ

LAMB

LAMB åĚňåijŔåęĆäÿŃïijŽ

$$m_t = \beta_1 m_{t-1} + (1 - \beta_1) g_t$$

$$v_t = \beta_2 v_{t-1} + (1 - \beta_2) g_t^2$$

$$r_t = \frac{m_t}{\sqrt{v_t} + \epsilon}$$

$$w_t = w_{t-1} - \eta_t \frac{\|w_{t-1}\|}{\|r_t + \lambda w_{t-1}\|}(r_t + \lambda w_{t-1})$$

åŠŇLARS çśżäijij, LAMB äź§æŸŕåIJÍåĘĚåśĆäijŸåŇŰåŹÍçŽĎå§žçąĂäÿŁ, åěŮäžĘäÿĂäÿłlocal learning rate çŽĎéĂżè¿Ś, åŕźæŕŔäÿĂåśĆçŽĎlearning rate è£ŻèąŇäžĘæŤ¿çijľãĂĆ

æŢĹædIJ

ä¡£çŤÍ LARS åIJÍèűĚåď ğbatch size äÿŃèő■çżČ resnet50ïijŽ

resnet50 imagenet	Global batch size	epoch	top1
[Goyal et al]	8k	90	76.3%
LARS Paper	32k	90	72.3%
[fleet: lars + amp]	16k	60	76.2%
[fleet: lars + amp]	32k	62	75.9%

ä¡£çŤÍæŰźæşŢ

LARS

fleet åřĘ LARSåőđçŐřäÿžäÿĂäÿł fleet meta optimizer, åIJÍä¡£çŤÍæŮűéIJĂèęĄèő¿ç¡őäžĕäÿŃåĞăçĆź:

1. LARS meta optimizer çŽĎ inner optimizer å£Ěéąžäÿž momentum, åźűåIJÍ momentum äÿ■åőŽäźL' mu åŠŇlr åŔĆæŢřãĂĆ
2. åIJÍDistributedStrategy äÿ■èő¿ç¡őLARS çL'źæIJL'çŽĎ lars_coeff åŔĆæŢřåŠŇ lars_weight_decay åŔĆæŢřãĂĆ
   - LARS åůšçżŔåřĘ weight decay åŇĚåŘńè£ŻåĚňåijŔäÿ■, çŤÍæĹůäÿ■éIJĂèęĄåĘ■åIJÍ optimizeräÿ■èő¿ç¡ő regularizationãĂĆ
   - fleet äÿ■è£ŸæŔŘä¿Ż lars_weight_decay è£ĞæžďçŰçŢě, åŘŕäžěéĂŽè£ĞåIJÍexclude_from_weight_decay åŔĆæŢřåĹååĚěåŕźåžŤlayer çŽĎ name string, èőľè£ŹäÿĂ layer çŽĎåŔĆæŢřäÿ■è£ŻèąŇ lars weight decayãĂĆ (éĂŽåÿÿæĹŚäžňåřĘBN åŔĆæŢř åŠŇFC_bias äžŐlars weight decay äÿ■è£Ğæžď')

```
strategy = fleet.DistributedStrategy()
strategy.lars = True
strategy.lars_configs = {
                    "lars_coeff": 0.001,
                    "lars_weight_decay": 0.0005,
                    "exclude_from_weight_decay": ['batch_norm', '.b_
↪0']
                }
```

äÿŁè£řä¿Ńå■ŘçŽĎåőŇæŢťäžčçăĄå■ŸæŤ¿åIJĺïijŽtrain_fleet_lars.pyäÿŃéİćãĂĆåĄĞèő¿èęĄè£ŘèąŇ2å■

```
fleetrun --gpus=0,1 train_fleet_lars.py
```

æĆĺåřĘçIJŃåĹřæŸ¿çd'žåęĆäÿŃæŮěå£Ůä£ąæĄŕïijŽ

```
----------  Configuration Arguments ----------
gpus: 0,1
heter_worker_num: None
heter_workers:
http_port: None
ips: 127.0.0.1
log_dir: log
...
-------------------------------------------------
...
+=================================================================================
|                          Distributed Envs                                       ␣
↪Value                      |
+--------------------------------------------------------------------------------
↪--------------------+
|                        PADDLE_TRAINER_ID                                  0 ␣
↪                    |
|                  PADDLE_CURRENT_ENDPOINT                  127.0.0.
↪1:12464               |
|                      PADDLE_TRAINERS_NUM                                2 ␣
↪                    |
|                 PADDLE_TRAINER_ENDPOINTS         127.0.0.1:12464,
↪127.0.0.1:43227       |
|                      FLAGS_selected_gpus                                0 ␣
↪                    |
+=================================================================================
...
+=================================================================================+
|                                                                                 ␣
↪            |
|                           DistributedStrategy Overview                          ␣
↪            |
|                                                                                 ␣
↪            |
+=================================================================================+
```

(äÿŃéąçżğçż■)

(çż■äÿŁéąţ)

```
|                          lars=True <-> lars_configs                          ␣
↪            |
+-----------------------------------------------------------------------------
↪-----------+
|                             lars_coeff          0.
↪0010000000474974513          |
|                       lars_weight_decay         0.
↪0005000000237487257          |
|                               epsilon                    0.0                ␣
↪            |
|              exclude_from_weight_decay                batch_norm             ␣
↪            |
|                                                          .b_0                ␣
↪            |
+=============================================================================+
...
W0114 18:07:51.588716 16234 device_context.cc:346] Please NOTE:␣
↪device: 4, GPU Compute Capability: 7.0, Driver API Version: 11.0,␣
↪Runtime API Version: 10.0
W0114 18:07:51.593963 16234 device_context.cc:356] device: 4, cuDNN␣
↪Version: 7.6.
[Epoch 0, batch 0] loss: 0.14651, acc1: 0.00000, acc5: 0.00000
[Epoch 0, batch 5] loss: 1.82926, acc1: 0.00000, acc5: 0.00000
[Epoch 0, batch 10] loss: 0.00000, acc1: 0.00000, acc5: 0.00000
[Epoch 0, batch 15] loss: 0.13787, acc1: 0.03125, acc5: 0.03125
[Epoch 0, batch 20] loss: 0.12400, acc1: 0.03125, acc5: 0.06250
[Epoch 0, batch 25] loss: 0.17749, acc1: 0.00000, acc5: 0.00000
...
```

åőŇæŢť 2å■ąçŽĎæŮěå£Ůä£ąæĄŕäź§åŔŕåIJĺ. `./log/çŻőå¡ŢäÿŃæ§ěçIJŃãĂĆäžĘèğčæŻť`åď`Žfleetr åŘŕåŁĺåĹĘåÿČåijŔäżżåŁąãĂĆ

LAMB

fleet åřĘ LAMBåőđçŐřäÿžäÿĂäÿł fleet meta optimizer, åIJĺä¡£çŤĺæŮűéIJĂèęĄèő¿ç¡őäžěäÿŃåĞăçĆź:

1. LAMB meta optimizer çŽĎ inner optimizer å£Ěéąžäÿž `adam`, åźűåIJĺ adam äÿ■åőŽäźL' å■ęäźăçŐĞ`lr`, äÿĂéŸű moment çŽĎæŇĞæŢřèąřåĞŔçŐĞ`beta1` åŠŇäžŇéŸűmoment çŽĎæŇĞæŢřèąřåĞŔçŐĞ`beta2` åŔĆæŢřãĂĆ
2. åIJĺ DistributedStrategy éĞŇåőŽèő¿ç¡őAMB çL'źæIJL'çŽĎ `lamb_weight_decay` åŔĆæŢř.
    - LAMB åůšçżŔåřĘ `weight decay` åŇĚåŘńè£ŻåĚňåijŔäÿ■, çŤĺæĹůäÿ■éIJĂèęĄåĘ■åIJĺ optimizeräÿ■èő¿ç¡ő `regularization`ãĂĆ
    - fleet äÿ■è£ŸæŔŘä¿Ż lamb_weight_decay è£Ğæžd'ç■ŰçŢě, åŕŕäžěéĂŽè£ĞåIJĺ`exclude_from_weight_decay`

åŔĆæŢřåŁäåĚěåŕźåžŤlayer çŽĎ name string, èőľè£ŹäÿĂ layer çŽĎåŔĆæŢřäÿ■è£ŻèąŇ lars weight decayãĂĆ (éĂŽåÿÿæĹŚäżňåřĘLN äżŐlamb weight decay äÿ■è£Ğæżďŕŕŕ)

```
strategy = fleet.DistributedStrategy()
strategy.lamb = True
strategy.lamb_configs = {
    'lamb_weight_decay': 0.01,
    'exclude_from_weight_decay': ['layer_norm'],
}
```

äÿŁè£řä¿Ńå■ŘçŽĎåőŇæŢťäžčçăĄå■ŸæŤ¿åIJĺïijŽtrain_fleet_lamb.pyäÿŃéİćãĂĆåĄĞèő¿èęĄè£ŘèąŇ2å

```
fleetrun --gpus=0,1 train_fleet_lamb.py
```

æĆĺåřĘçIJŃåĹřæŸ¿çď˙žåęĆäÿŇæŮěå£Ůä£ąæĄŕïijŽ

```
-----------  Configuration Arguments -----------
gpus: 0,1
heter_worker_num: None
heter_workers:
http_port: None
ips: 127.0.0.1
log_dir: log
...
------------------------------------------------
...
+=================================================================================
|                          Distributed Envs                                      ␣
↪Value                       |
+--------------------------------------------------------------------------------
↪-------------------+
|                        PADDLE_TRAINER_ID                                    0 ␣
↪                    |
|                  PADDLE_CURRENT_ENDPOINT                  127.0.0.
↪1:12464               |
|                      PADDLE_TRAINERS_NUM                                  2 ␣
↪                    |
|                  PADDLE_TRAINER_ENDPOINTS         127.0.0.1:12464,
↪127.0.0.1:43227       |
|                      FLAGS_selected_gpus                                  0 ␣
↪                    |
+=================================================================================
...
+=================================================================================+
|                                                                                ␣
↪            |
|                           DistributedStrategy Overview                         ␣
↪            |
|                                                                                ␣
↪            |
```

(äÿŃéąçżğçż■)

(çż■äÿŁéąţ)

```
+=====================================================================+
|                          lamb=True <-> lamb_configs                 ␣
↪           |
+-----------------------------------------------------------------
↪-----------+
|                     lamb_weight_decay           0.
↪009999999776482582          |
|              exclude_from_weight_decay                layer_norm   ␣
↪           |
+=====================================================================+
...
W0114 18:07:51.588716 16234 device_context.cc:346] Please NOTE:␣
↪device: 4, GPU Compute Capability: 7.0, Driver API Version: 11.0,␣
↪Runtime API Version: 10.0
W0114 18:07:51.593963 16234 device_context.cc:356] device: 4, cuDNN␣
↪Version: 7.6.
[Epoch 0, batch 0] loss: 0.14651, acc1: 0.00000, acc5: 0.00000
[Epoch 0, batch 5] loss: 1.82926, acc1: 0.00000, acc5: 0.00000
[Epoch 0, batch 10] loss: 0.00000, acc1: 0.00000, acc5: 0.00000
[Epoch 0, batch 15] loss: 0.13787, acc1: 0.03125, acc5: 0.03125
[Epoch 0, batch 20] loss: 0.12400, acc1: 0.03125, acc5: 0.06250
[Epoch 0, batch 25] loss: 0.17749, acc1: 0.00000, acc5: 0.00000
...
```

åőŇæŢť2 å■ąçŽĎæŮěå£Ůä£qæĄŕäź§åŔŕåIJÍ./log/çŻőå¡ŢäÿŃæ§ěçIJŃãĂĆäžĘèğčæŻť åď ŽfleetråŘŕåŁĺåĹĘåÿČåijŔäżżåŁąãĂĆ

6.5.4 éčđæąĺåď ğèğĎæĺąåĹĘçśżåžŞä¡£çŤĺäżŃçż■

çőĂäżŃ

åŻ¿åČŔåĹĘçśżæŁĂæIJŕæŮěèűŃæĹŘçĘ§ïijŇResNetç¡ŚçżIJåIJÍImageNetæŢřæ■őéŻĘäÿŁçŽĎtop5åĞĘç

äżŐåď ŽåĹĘçśżçěđçżŔç¡ŚçżIJçŽĎåőđçŐřèğŠåžęåĹĘæđŘïijŇåĚűæIJĂåŘŐäÿĂåśĆéĂŽåÿÿæŸŕçŤśåĚĺè£

äžěæŰřéŮżæŐĺè■ŘçşżçżŞäÿžä¿ŃïijŇåĄĞèő¿èęĄåŕźçŹ¿äÿĞçśżçŻĘåĹĘçśżåĹńçŽĎæŰřéŮżæİąçŽőè£Żè

$$åĚĺè£đæŐěåśĆåŔĆæŢřæŸ¿å■ŸæűĹèĂŮ=\frac{512*10^6*4B}{1024^3}\approx2GB$$

åŐ§çŘĘäżŃçż■

èŕěåęĆä¡Ţèğčåę£è£ŹäÿłéŮőéćŸåŚćïij§åÿÿçŤĺçŽĎåĄŽæşŢæŸŕâĂIJæŃĘåĹĘâĂİãĂĆèĂČèŹŚåĹřåĚĺè£đ

äžěäÿŃåŻ¿äÿžä¿ŃïijŇåĚĺè£đæŐěåśĆåŔĆæŢřæŇĽèąŇåĽĞåĹĘåĹřäÿ■åŘŇçŽĎGPUå■ąäÿŁãĂĆærŔæň
logit)ïijŇåźűå§žäžŐæ■ď èőąçőŮçěđçżŔç¡ŚçżIJçŽĎæ■§åď śåĂijãĂĆèŕęçżĘæŐĺåŕijè£ĞçĺŃèŕůåŔĆéŸĚéŹĎå¡

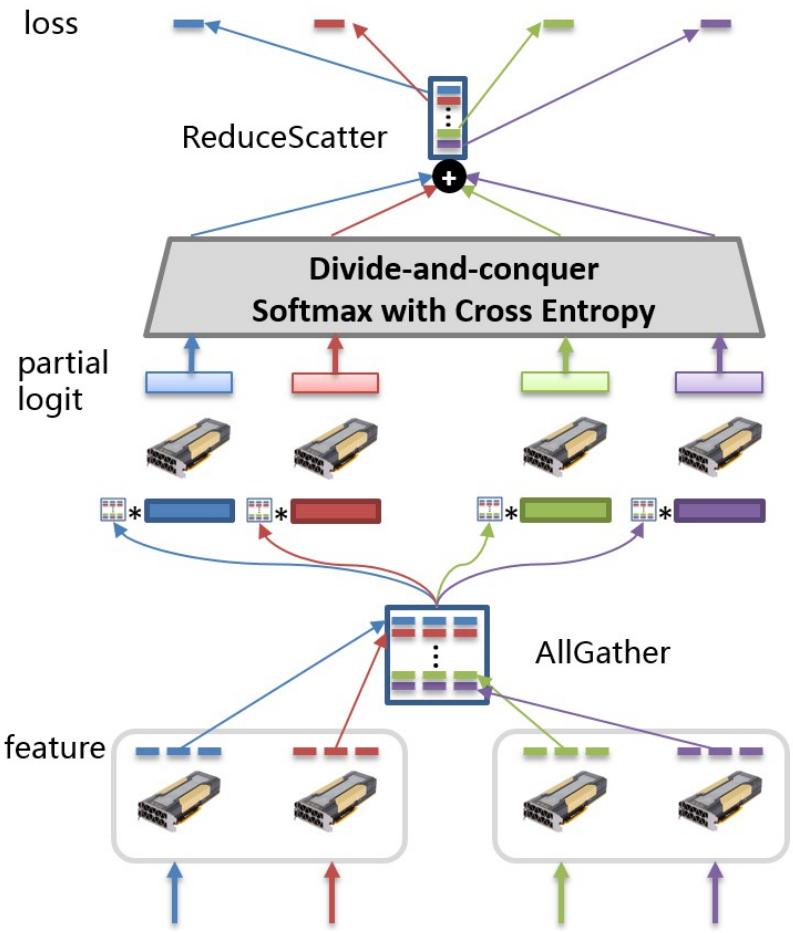

è£ŹäÿłæŰźæąĹåŕŕäžěæIJL'æŢĹèğčåĘşåĚĺè£đæŐěåśĆåŔĆæŢřéĞŔéŽŔåĹĘçśżçśżåĹńæŢřçž£æĂğåćđéŢ
PaddlePaddle Large Scale Classification)ïijŇäÿžçŤĺæĹůæŔŘä¿ŻäžĘåd'ğèğĎæĺąåĹĘçśżäżżåŁąäżŐèő■çżČåĹř

åŁ§èČ¡æŢĹæđlJ

PLSCåžŞåIJĺåd'ŽäÿłæŢřæ■őéŻĘäÿŁåŔŕäžěåŔŰå¿ŮSOTAçŽĎèő■çżČçš¿åžęïijŇäÿŃèąĺåĹŮåĞžPLSCåž
ArcFaceåŠŇCASIAæŢřæ■őéŻĘä¡IJäÿžèő■çżČæŢřæ■őïijŇåIJĺäÿ■åŘŇéłŇèŕĄæŢřæ■őéŻĘäÿŁåŔŰå¿ŮçŽĎçš

æĺąåđŃ	èő■çżČéŻĘ	lfw	agendb_30	cfp_ff	cfp_fp	MegaFace (Id/Ver)
ResNet50	MS1M-ArcFace	0.99817	0.99827	0.99857	0.96314	0.980/0.993
ResNet50	CASIA	0.98950	0.90950	0.99057	0.91500	N/A

åd'ĞæşĺïijŽäÿŁè£řæĺąåđŃèő■çżČä¡£çŤĺçŽĎloss_typeäÿžâĂŹdist_arcfaceâĂŹãĂĆæŻt'åd'ŽåĚşäžŐArcF
ArcFace: Additive Angular Margin Loss for Deep Face RecognitionãĂĆ

PLSCæŤŕæŇĄåd'ŽæIJžåĹĘåÿČåijŔèő■çżČãĂĆäÿĂæŰźéĺćïijŇéĂŽè£Ğåd'ŽæIJžåĹĘåÿČåijŔèő■çżČåŔ
GPUéĚ■ç¡őäÿŃæŤŕæŇĄçŽĎæIJĂåd'ğåĹĘçśżçśżåĹńæŢřçŻÿæŕŤäÿ■ä¡£çŤĺPLSCæL'l'åd'ğ2.52åĂ■ãĂĆåŔęäÿ

äÿŃåŻ¿çżŹåĞžä¡£çŤĺäÿ■åŘŇæŢřéĞŔçŽĎèŁĆçĆźæŰűçŽĎèő■çżČéĂ§åžęïijĹåŘđåŘŘïijL'ãĂĆåőđéłŇä
ArcFaceïijŇåĹĘçśżçśżåĹńæŢřäÿž85742ïijŇæŕŘäÿłèŁĆçĆźéĚ■åd'Ğ8åijãNVIDIA V100
GPUsïijŇbackboneæĺąåđŃäÿžResNet50ãĂĆåęĆåŻ¿æL'Äçd'žïijŇä¡£çŤĺéčđæąĺåd'ğèğĎæĺąåĹĘçśżåžŞåŕŕäžěă

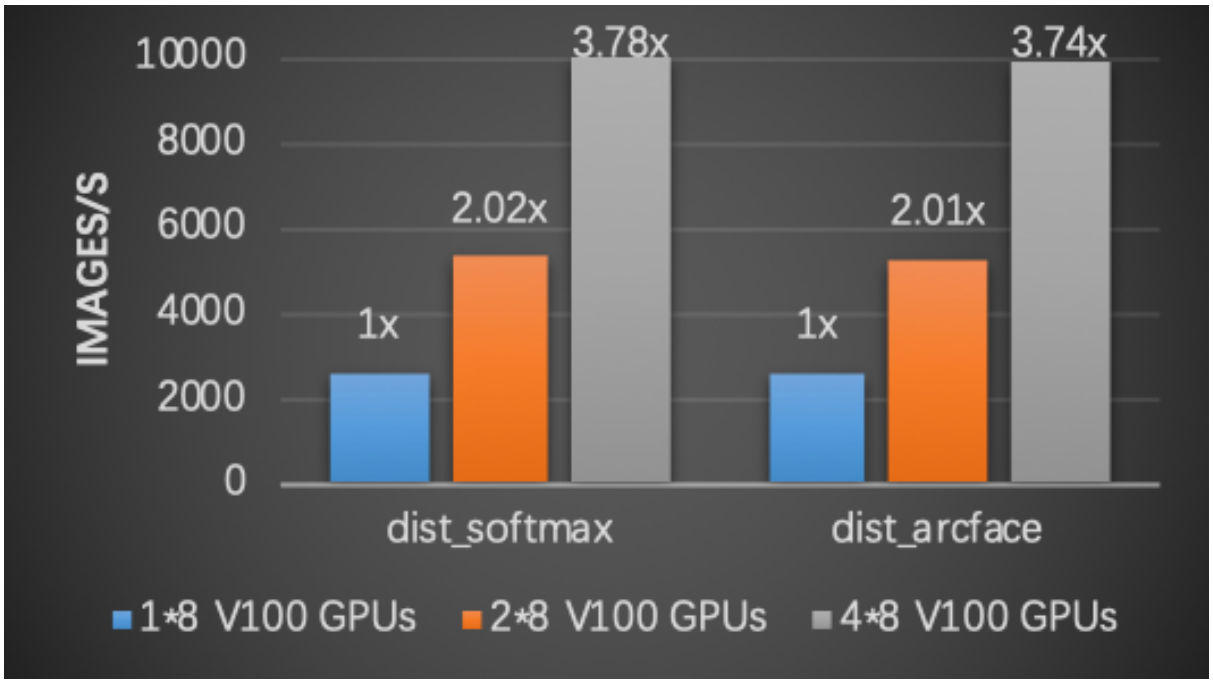

ä¡£çŤÍæŰźæşŢ

åőĽèčĚPLSC

æĽğèąŇäÿŃéİćçŽĎåŚ¡äžďåőĽèčĚPLSCãĂĆ

```
pip install plsc
```

åĞĘåďĞæÍąåđŃèő■çżČéĚ■ç¡őäžčçăĄïijŇä£Ìå■Ÿäÿžtrain.pyæŰĞäžű

ä¡£çŤÍPLSCçżĎåžžåĹĘçśžçěđçżŔç¡ŚçżIJäÿżèęĄåŇĚæŃňäÿŃéİćäÿĽäÿłæ■ěéłď¨ijŽ

1. äžŐplscåŇĚåŕijåĚěEntryçśžïijŇEntryçśžåřĄèčĚPLSCæĽĂæIJĽAPIçŽĎæŐěåŔčçśžïijŻ
2. åőđä¿ŃåŇŰEntryçśžçŽĎåŕźèśąïijŻ
3. èřČçŤÍEntryçśžçŽĎtrainæŰźæşŢïijŇåijĂåğŃèő■çżČè£ĞçÍÑãĂĆ

éżŸèőď¨æČĚåĘţäÿŃïijŇèřěèő■çżČèĎŽæIJñä¡£çŤÍçŽĎlossåĂijèőąçőŮæŰźæşŢäÿžâĂŹdist_arcfaceâĂŹiæŰĞæąčãĂĆ

```
from plsc import Entry
if __name__ == "main":
        ins = Entry()
        ins.set_class_num(1000000) #èő¿ç¡őåĹĘçśžçśžåĹńæŢř
        ins.train()
```

åŘŕåŁĺèő■çżČäżżåŁą

åŘŕäžěä¡£çŤĺäÿŃéĺćçŽĎåŚ¡äžd'èąŇåŘŕåŁĺèő■çżČäżżåŁąïijŇåĚűäÿ■selected_gpusåŘĆæŢřçŤĺäžŐæŇĞ

```
python -m paddle.distributed.launch \
            --selected_gpus=0,1,2,3,4,5,6,7 \
            train.py
```

æŻt'åd'ŽPLSCä¡£çŤĺæŰĞæąčïijŇèŕůåŘĆéŸĚ: PLSC RepoãĂĆ

éŹĎå¡Ţ

åĚĺè£đæŐěåśĆæŞ■ä¡IJåIJĺæŢřå■ęäÿŁç■L'äžůäžŐè¿ŞåĚěXåŠŇåŘĆæŢřWçŽĎç§l'éŸţäźŸ:
XWãĂĆåŘĆæŢřWåŘŕäžěæŇL'åĹŮåĹĞåĹĘäÿžNäÿłéČĺåĹĘ $[W_0, W_1, ..., W_{N-1}]$ïijŇåźűåĹĘåĹńæŤ¿ç¡őåĹřN

$$XW = X[W_{0}, W_{1}, âĂę, W_{N-1}] = [XW_{0}, XW_{1}, âĂę, XW_{N-1}]$$

åŻăæ■d'ïijŇåIJĺçňñiåijåå■ąäÿŁïijŇåŘłéIJĂèęĄèőąçőŮéČĺåĹĘçżŞædIJ
XW_iãĂĆçĎűåŘŐïijŇéĂŽè£ĞéŽĘåŘĹéĂŽä£ąæŞ■ä¡IJèŐůåŘŮåĚĺåśĂçżŞædIJ XWãĂĆ

6.5.5 ä¡£çŤĺSharding èő■çżČèűĚåd'ğæĺąåđŃ

çőĂäżŃ

å¡ŞæĺąåđŃåŘĆæŢřè¿¿åĹřçŹ¿äž£æĹŰèĂĚå■Čäž£æŮűïijŇ
äijăçż§çŽĎæŢřæ■őåźűèąŇèő■çżČåŕřèČ¡äijŽéĄĞåĹřæŸ¿å■ŸçŞűécĹãĂĆ
åIJĺæŢřæ■őåźűèąŇèő■çżČäÿ■ïijŇæŕŘäÿłgpu worker éČ¡æIJL'äÿĂäż¡åőŇæŢt'æĺąåđŃåŘĆæŢřåŠŇäijŸåŇŰåŹ
ãĂŁZeRO: Memory Optimizations Toward Training Trillion Parameter ModelsãĂŃ
æŇĞåĞžåIJĺæŕŘäÿłGPU äÿŁéČ¡ä£Ìå■ŸäÿĂäż¡æĺąåđŃåŘĆæŢřåŠŇäijŸåŇŰåŹĺçŁűæĂĄåL'ŕæIJňæŸŕåĘŮä¡Z
æĹŚäžňåŘŕäžěéĂŽè£ĞářĘäÿŁè£řåŘĆæŢřåŠŇåL'ŕæIJňåĹŚåĹĘåĹřäÿ■åŘŇGPU äÿ■ïijŇ
åIJĺæŕŘäÿłGPU åŘłä£Ìå■ŸéČĺåĹĘåL'ŕæIJňïijŇæİěåĞŘářŚæŕŘåijăGPUäÿŁæŸ¿å■ŸçŽĎå■ăçŤĺïijŇäžŐèĂŇål

åŐ§çŘĘ

Sharding

Sharding åőđçŐřäžEçśżäijijZeRO-DP çŽĎèő■çżČç■ŰçŢěïijŇářĘæĺąåđŃåŘĆæŢřïijĹpa-
rameterïijL'ãĂĄåŘĆæŢřæćŕåžęïijĹgradientïijL'ãĂĄåŘĆæŢřåŕźåžŤçŽĎäijŸåŇŰåŹĺçŁűæĂĄïijĹmomentïijL'ä
äÿŁãĂČèől'æĺąåđŃåŘĆæŢřéČĺåĹĘæL'Ăå■ăçŽĎæŸ¿å■ŸéŽŘåźűèąŇå■ąæŢřçŽĎåćđåŁăèĂŇåĞŘářŠãĂĆ
éĂŽè£Ğ paddle.distributed.fleet æŘŘä¿ŻçŽĎçőĂå■ŢæŸŞçŤĺæŐěåŘč,
çŤĺæĹůåŘłéIJĂèęĄæůżåŁăåĞăèąŇäžčçăĄåřśåŘŕářEç■ŰçŢěåŁăåĚěåĹřåŐ§æIJL'çŽĎèő■çżČäÿ■ãĂĆ

æĺąåđŃèő■çżČè£ĞçĺŃäÿ■çŽĎæŸ¿å■Ÿæűĺè ĂŮäÿżèęĄçŤśäÿd'åd'ğéČĺåĹĘçżĎæĹŘïijŽæĺąåđŃåŘĆæŢř
sharding ç■ŰçŢěåŘłåĹĞåĹĘäžĘæĺąåđŃåŘĆæŢřåŠŇäijŸåŇŰåŹĺçŁűæĂĄïijŇåŽăæ■d'æĺąåđŃåŘĆæŢřåŠŇäij
ä¡ĘæŸŕæŕŘäijăGPUäÿŁäż■çĎűçżt'æŁd'çİĂæĺąåđŃåőŇæŢt'çŽĎåL'■åŘŚåŠŇåŘ■åŘŚïijŇæL'ĂäżěæŕŘäijăGP

äÿ■äijŽéŽŕçİĂGPU æŢřéĞŕçŽĎåćđåŁäèĂŇåĞŕåřŚãĂĆ çŤĺæĹůåŕŕäžěéĂŽè£ĞçżŞåŘĹ recompute ç■ŰçŢĕæİěåĞŕåřŚ activationè£ŹéČĺåĹEçŽĎæŸ¿å■ŸæűĹèĂŮãĂĆ

éĂŽè£Ğsharding åŠŇåćđåŁäåźűèąŇGPU æŢřéĞŕïijŇçŤĺæĹůåŕŕäžěèő■çżČçŹ¿äž£çŤŽèĞşå■Čäž£åŕĆé ïijĹéIJĂèęĄçżŞåŘĹ recompute, amp ç■ŰçŢěïijL'ãĂĆ

Sharding-hybrid-dp

Sharding hybridæŢřæ■őåźűèąŇç■ŰçŢěïijŇåIJĺsharding åźűèąŇçŽĎå§žçąĂäÿŁåĘ■åćđåŁääÿĂåśĆæŢřæ■ èŕěç■ŰçŢěçŽĎçŻőçŽĎæŸŕéĂŽè£Ğ éŹŘåĹűsharding éĂŽä£ąçŽĎèŁĆçĆźæŢř åŠŇ åćđåŁäåd′ŽèůŕæŢřæ■őåźűèąŇ æİěæŔŘéńŸèő■çżČåŘđåŘŘãĂĆ åęĆæđIJäÿĂäÿłæĺąåđŃåIJĺæŹőéĂŽSharding èő■çżČæŮűéIJĂèęĄM åijäGPUïijŇåĹŹåĹŹåijĂåŘŕhybrid-dp èĞşåřŚéIJĂèęĄ N*M GPU ïijĹN>= 2ïijL'ãĂĆ

Sharding-hybrid-dp éĂĆçŤĺçŽĎåIJžæŹŕåęĆäÿŃïijŽ

- å¡ŞåL'■æIJL' 4äÿł 8 å■ąv100 èŁĆçĆź
- çŻőæăĞæĺąåđŃA åIJĺSharding èő■çżČæŮűèĞşåřŚéIJĂèęĄ 8å■ą v100 ïijĹäÿĂäÿłåőŇæŢ'çŽĎ8 å■ąv100èŁĆçĆźïijL'
- åÿŇæIJŻåĹl'çŤĺåĚĺéČĺçŽĎ 4 äÿłèŁĆçĆźæİěåŁäéĂ§èő■çżČ

äÿŁè£řæČĚåEţåęĆæđIJçŻt'æŐëä¡£çŤĺåĚĺéČĺçŽĎ 4 äÿłèŁĆçĆź è£ŻèąŇæŹőéĂŽçŽĎsharding èő■çżČïijŇ éĆčäźĹåĚĺéČĺçŽĎ 32 gpus äźŃéŮt'çżĎæĹŘäÿĂäÿłåőŇæŢt' Sharding parallelismãĂĆè£ŹæăůäijŽåŻääÿžéĂŽä£ąçŞűécĹéĂäæĹŘèő■çżČéĂ§åžęéİđåÿÿæĚćïijŽ

- Sharding äÿ■çŽĎbroadcast éĂŽä£ą äijŽæűL'åŘŁåĚĺéČĺçŽĎ32 åijăå■ąïijŇäÿŤäÿžèůĺèŁĆçĆźéĂŽä£ąãĂĆ
- Sharding äÿ■çŽĎ allreduce éĂŽä£ą äijŽæűL'åŘŁåĚĺéČĺçŽĎ32 åijăå■ąïijŇäÿŤäÿžèůĺèŁĆçĆźéĂŽä£ąãĂĆ

åijĂåŘŕ hybrid-dp åźűèő¿ç¡ő `sharding_group_size = 8` åŘŐïijŇ æŕŘäÿłèŁĆçĆźåĘĚçŽĎ 8 åijăå■ąçżĎæĹŘäÿĂäÿłåőŇæŢt'çŽĎ Sharding parallelismïijŇ4 äÿłèŁĆçĆźæđĎæĹŘ 4èůŕ hybrid data parallelismïijŽ

- Sharding äÿ■çŽĎbroadcast éĂŽä£ąèćńéŹŘåĹűåIJĺæŕŘäÿłèŁĆçĆźåĘĚçŽĎ 8 åijäGPU äźŃéŮt'ïijŇ æšąæIJL'èůĺèŁĆçĆźéĂŽä£ąãĂĆ
- Sharding äÿ■çŽĎ allreduce äÿžèůĺèŁĆçĆźéĂŽä£ąïijŇä¡ĘæŕŘäÿłallreduce éĂŽä£ąåŘłæűL'åŘŁ åŕźåžŤ sharding_group äÿŁ rank çŻÿåŘŇçŽĎ 4 åijäGPUsïijŇ äÿŤæŕŘåijäGPUäżĚéIJĂèęĄ allreduceéĂŽä£ą 1/8 çŽĎæĺąåđŃåŔĆæŢřãĂĆ

Sharding-hybrid-dp éĂŽè£ĞäÿŁè£řæŐłæŰ¡ïijŇåŘŕäžěè¿Čåd'ğçĺŃåžę åĞŘåřŚ Sharding èő■çżČ äżŐ1èŁĆçĆźæL'l'åśŢåĹř4 èŁĆçĆźæŮűçŽĎïijĹèůĺèŁĆçĆźïijL'éĂŽä£ąéĞŘãĂĆæŘŕéńŸèŁĆçĆźåćđå

P.S. hybrid dp æŸŕåŻääÿž Sharding parallelism æIJňèžńåĘĚåŘńäÿĂåśĆ data parallelism éĂżè¿ŚïijŇ hybrid dp æŸŕåIJĺ Sharding parallelismäźŃäÿŁåĘ■åćđåŁäæŰřçŽĎäÿĂåśĆ data parallelism éĂżè¿ŚãĂĆ

æŢĹædIJ

äÿŃéİćèąÍæäijåřĘåŕźæŕŤ Sharding ç■ŰçŢěåŕźæŸ¿å■ŸçŽĎå¡śåŞ■ãĂĆ

æÍąåđŃäÿž Bert-LargeïijŇèŕŢéłŇçŐŕåćČäÿž v100 ïijĹ32GBïijL'ïijŇ recompute = ON, amp = ON, hybrid-dp = OFFãĂĆ æÍąåđŃäÿ■åŔŸïijŇå■Ţå■ąbatch size äÿ■åŔŸïijŇå¡Şåźűèą ŇGPUæŢřéĞŔåćđåŁäæŮűïijŇæŸ¿å■ŸçŽĎæűĹèĂŮåřĘåĞŔåřŘãĂĆ çIJĄäÿŃçŽĎæŸ¿å■ŸåŔŕäžěçŤÍæİěåćđåď ğæÍąåđŃãĂĆ

setting	GPU Mem
shardingâĂŤoff	8667 MB
shardingâĂŤon N1C2	5615 MB
shardingâĂŤon N1C4	4841 MB
shardingâĂŤon N1C8	4127 MB
shardingâĂŤon N2C16	3700 MB

Sharding çżŞåŘĹ amp + recomputeïijŇåŔŕäžěåIJÍ 128 åijă 32GB V100 åźűèąŇçŽĎæČĚåĘţäÿŃæŤŕæŇĄå■Čäž£åŔĆæŢřïijĹ115BïijL'ERNIE èő■çżČãĂĆ

ä¡£çŤÍæŰźæşŢ

åŕźäžŐshardingïijŇçŤÍæĹůéIJĂèęĄèő¿ç¡ő `fuse_broadcast_MB`åŔĆæŢřãĂĆèŕěåŔĆæŢřæŐğåĹűåź£èő■çżČäÿ■çŽĎéĂŽä£ąéĂ§åžęïijŇæŸŕäÿĂäÿłéIJĂèęĄæäźæ■őåĚůä¡ŞæÍąåđŃåď ğåřŘåŠŇç¡ŚçżIJæŃŞæL'Śèő

èŃěåijĂåŘŕhybrid-dpïijŇçŤÍæĹůéIJĂèęĄèő¿ç¡ő `hybrid_dp`äÿžTrueïijŇåźűæŇĞåőŽ `sharding_group_size`ãĂĆ

äÿžäžĘçď žä¿Ńäžčçă ĄçŽĎçőĂçżČïijŇäÿŃéİćä¿Ńå■Řäÿ■ä¡£çŤÍè¿ČåřŘ resnet50 æÍąåđŃä¡IJäÿžæijŤçď žãĂĆåőđéŹĚèő■çżČäÿ■ïijŇsharding çŽĎçŽőæăĞæŸŕéĂŽè£ĞçL'žçL'šèő■çżČéĂ§åžęäžěæ■ćåŔŮåŕźæŻť åď ğæÍąåđŃçŽĎæŤŕæŇĄïijŇæŢĚäÿ■éĂĆ resnet50 ç■L'å■Ţå■ąåřśèČ¡èő■çżČçŽĎæÍąåđŃãĂĆ

åŻääÿžresnet50 è¿ČåřŔïijŇæĹŚäžňåŔŕäžěäžď `sharding_group_size = 2`èőľ æÍąåđŃåŔĆæŢřèćńåĹĞåĹĘäÿž2 äÿłshardsïijŇçĎůåŘŐåIJÍ äÿĂäÿłå■ŢæIJž4å■ą v100 çŽĎèŁĆçĆźäÿŁçżĎæĹŘ 2 èůŕ hybrid-dp åźűèąŃè£ŻèąŇæijŤçď žãĂĆ

```
strategy = fleet.DistributedStrategy()
strategy.sharding = True
strategy.sharding_configs = {
    "fuse_broadcast_MB": 32,
    "hybrid_dp": True,
    "sharding_group_size": 2,
}
```

äÿŁè£ŕä¿Ńå■ŘçŽĎåőŇæŢť äžčçăĄå■ŸæŤ¿åIJĺïijŽtrain_fleet_sharding.pyäÿŃéİćãĂĆåĄĞèő¿èęĄè£Řèąŕ

```
fleetrun --gpus=4,5,6,7 train_fleet_sharding.py
```

æĆĺåřĘçIJŃåĹřæŸ¿çď žåęĆäÿŇæŮěå£Ůä£ąæĄŕïijŽ

```
----------  Configuration Arguments ----------
gpus: 4,5,6,7
heter_worker_num: None
heter_workers:
http_port: None
ips: 127.0.0.1
log_dir: log
...
-----------------------------------------------
...
+=======================================================================================
|                        Distributed Envs                                 ␣
 ↪Value                      |
+---------------------------------------------------------------------
 ↪--------------------+
|                        PADDLE_TRAINER_ID                           0 ␣
 ↪                    |
|                  PADDLE_CURRENT_ENDPOINT                  127.0.0.
 ↪1:18362               |
|                      PADDLE_TRAINERS_NUM                           4 ␣
 ↪                    |
|                 PADDLE_TRAINER_ENDPOINTS  ... 0.1:23911,127.0.0.
 ↪1:35135,127.0.0.1:38263|
|                      FLAGS_selected_gpus                           4 ␣
 ↪                    |
+=======================================================================================
...
INFO:root:Using Sharing&DP mode !
INFO:root:global word size: 4
INFO:root:global rank: 0
INFO:root:sharding group_size: 2
INFO:root:sharding rank: 0
INFO:root:dp group size: 2
INFO:root:dp rank: 0
INFO:root:current endpoint: 127.0.0.1:18362
INFO:root:sharding group endpoints: ['127.0.0.1:18362', '127.0.0.
 ↪1:23911']
INFO:root:dp group endpoints: ['127.0.0.1:18362', '127.0.0.1:35135']
INFO:root:global word endpoints: ['127.0.0.1:18362', '127.0.0.
 ↪1:23911', '127.0.0.1:35135', '127.0.0.1:38263']
server not ready, wait 3 sec to retry...
not ready endpoints:['127.0.0.1:23911']
...
+=======================================================================================+
|                       sharding=True <-> sharding_configs                ␣
 ↪           |
+---------------------------------------------------------------------
 ↪-----------+
|                      fuse_broadcast_MB                     32.0      ␣
 ↪           |
```

(äÿŃéąţçżğçż■)

(çż■äÿŁéąţ)

```
|                                 hybrid_dp                                 True      ␣
↪            |
|                      sharding_group_size                                   2         ␣
↪            |
+=======================================================================================+
...
W0114 18:07:51.588716 16234 device_context.cc:346] Please NOTE:␣
↪device: 4, GPU Compute Capability: 7.0, Driver API Version: 11.0,␣
↪Runtime API Version: 10.0
W0114 18:07:51.593963 16234 device_context.cc:356] device: 4, cuDNN␣
↪Version: 7.6.
[Epoch 0, batch 0] loss: 0.14651, acc1: 0.00000, acc5: 0.00000
[Epoch 0, batch 5] loss: 1.82926, acc1: 0.00000, acc5: 0.00000
[Epoch 0, batch 10] loss: 0.00000, acc1: 0.00000, acc5: 0.00000
[Epoch 0, batch 15] loss: 0.13787, acc1: 0.03125, acc5: 0.03125
[Epoch 0, batch 20] loss: 0.12400, acc1: 0.03125, acc5: 0.06250
[Epoch 0, batch 25] loss: 0.17749, acc1: 0.00000, acc5: 0.00000
...
```

åőŇæŢť4å■ąçŽĎæŮěå£Ůä£ąæĄŕäź§åŔŕåIJÍ. `/log/çŽőå¡ŢäÿŃæ§ěçIJŃãĂĆäžĘèğčæŻťåď Žfleetru åŘŕåŁÍåĹĘåÿČåijŔäżżåŁąãĂĆ`

6.5.6 æţĄæřťçž£åźűèąŇ

çőĂäżŃ

éĂŽåÿÿæÌěèőšïijŇèő■çżČæŻťåď ğèğĎæÍąçŽĎç¡ŚçżIJæÍąåđŃåŔŕäžěåIJÍåď Žçğ■äżżåŁąäÿŁåŔŰå¿ŮæŻť

åŐ§çŘĘäżŃçż■

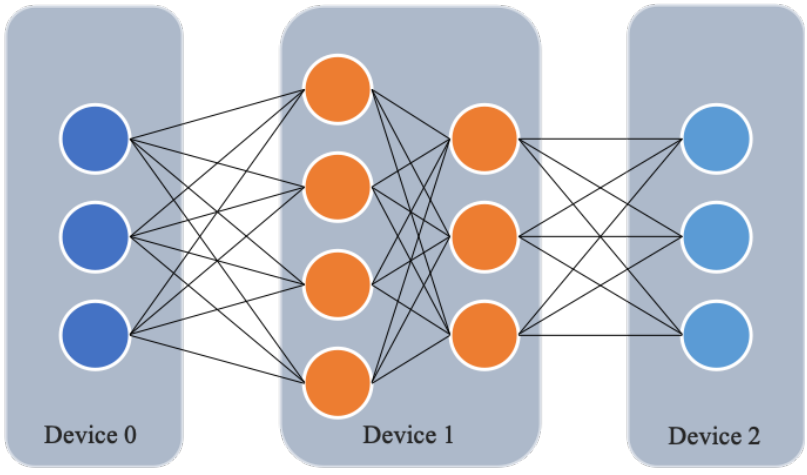

äÿŐæŢřæ■őåźűèąŇäÿ■åŘŇïijŇæţĄæřťçž£åźűèąŇåřĘæÍąåđŃçŽĎäÿ■åŘŇåśĆæŤ¿ç¡őåĹřäÿ■åŘŇçŽĎèő

äÿŃåŻ¿çżŹåĞžæţĄæřťçž£åźűèąŇçŽĎæŮűåžŔåŻ¿ãĂĆæIJĂçőĂéĚ■ç¡őæţĄæřťçž£åźűèąŇæÍąåđŃäÿŃïij

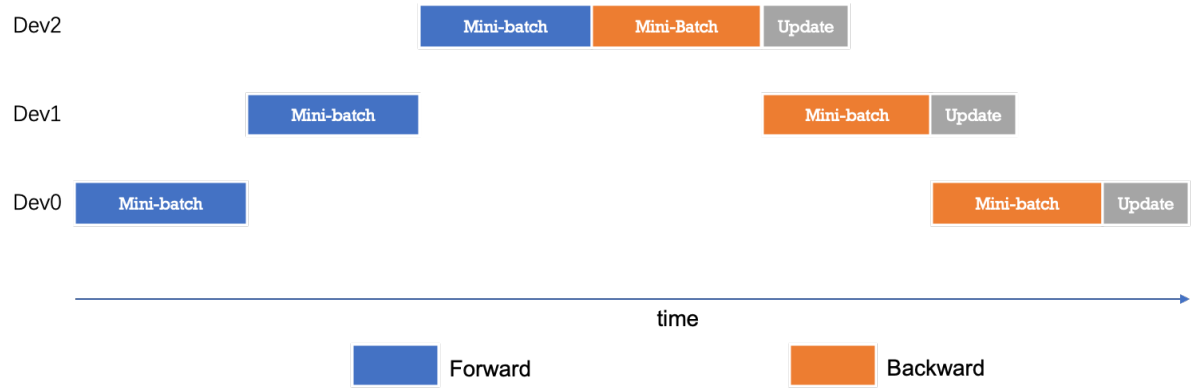

äÿžäžĘäijŸåŇŰæţĄæřť çž£åźűèąŇäÿ■èő¿åď ĞçŽĎèőąçőŮæŢĹçŐĞïijŇåŕŕäžěè£ŻäÿĂæ■ěårĘmini-batchåĹĞåĹĘæĹŘèŇěåźšæŻť årŔçšŠåžęçŽĎmicro-batchïijŇäžěæŔŘå■ĞæţĄæřť çž£åźűèąŇçŽĎåźűåŔŠåžęïijbatchèćńåĹĞåĹĘäÿž4äÿłmicro-batchïijŻåL’■åŘŠéŸűæőţïijŇæŕŔäÿłèő¿åď Ğä¿Ìæňąèőąçőůå■Ţäÿłmicro-batchçŽĎçżŞæđIJïijŻäžŐèĂŇåćđåŁääžĘèő¿åď ĞéŮť çŽĎåźűåŔŠåžęïijŇéŹ■ä¡ŐäžĘæţĄæřť çž£åźűèąŇbubbl

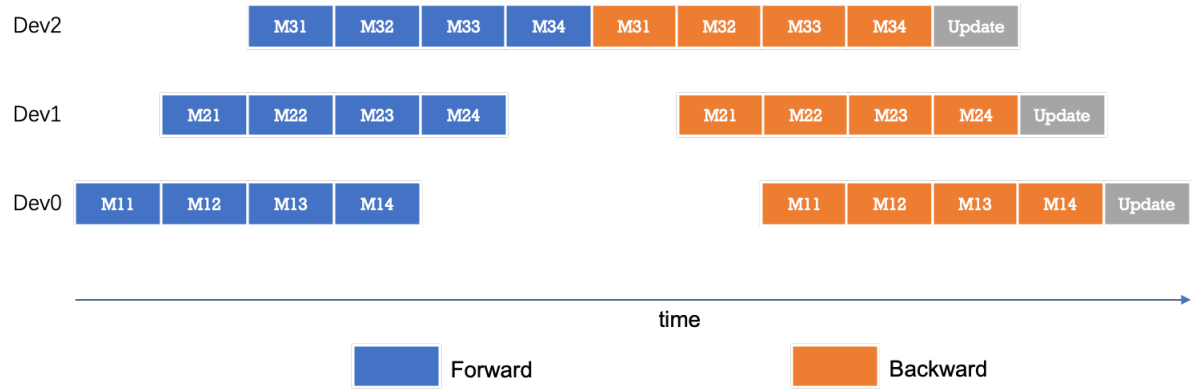

åŁ§èČ¡æŢĹædIJ

ä¡£çŤĺæţĄæřť çž£åźűèąŇïijŇåŕŕäžěåőđçŐřèűĚåď ğèğĎæĺąæĺąåđŃèő■çżČãĂĆä¿ŃåęĆïijŇä¡£çŤĺåď Žäÿ

ä¡£çŤĺæŰźæşŢ

åIJĺä¡£çŤĺæţĄæřť çž£åźűèąŇçŽĎèő■çżČç■ŰçŢěæŮűïijŇæĹŚäžňéĂŽè£ĞdevicedæŐěåŔčårĘ

```
# device_guard ä¡£çŤĺçd’žä¿Ń
def build_network():
    with paddle.fluid.device_guard("gpu:0"):
        data = paddle.static.data(name='sequence', shape=[1], dtype=
↪'int64')
        data_loader = paddle.io.DataLoader.from_generator(
            feed_list=[data],
            capacity=64,
            use_double_buffer=True,
            iterable=False)
        emb = nn.embedding(input=data, size=[128, 64])
    with paddle.fluid.device_guard("gpu:1"):
```

(äÿŃéąţçżğçż■)

(çż■äÿŁéąţ)

```
        fc = nn.fc(emb, size=10)
        loss = paddle.mean(fc)
    return data_loader, loss
```

éĂŽè£Ğèő¿åőŽdist_strategy.pipeline äÿžTrueïijŇåřĘæţĄæřť çž£åźűèąŇçŽĎç■ŰçŢěæ£Ăæť

```
fleet.init(is_collective=True)
dist_strategy = paddle.distributed.fleet.DistributedStrategy()
dist_strategy.pipeline = True
```

è£ŽäÿĂæ■ěåIJřïijŇåŕŕäžěéĂŽè£Ğdist_strategy.pipeline_configs éĚ■ç¡őæţĄæřť çž£åźűèąŇäÿ■mini-batchçŽĎåĹĞåĹEçšŠåžęãĂĆåĄĞèő¿mini-batchçŽĎåď ğåřŕäÿž128ïijŇåŕŕäžěéĂŽè£ĞäÿŃè£řäžčçăĄåřĘmini-batchåĹĞäÿž4äż¡æŽť åřŔçšŠåžęçŽĎmicro-batchïijŇæŕŕäÿłmicro-batchçŽĎåď ğåřŕäÿž32ãĂĆéIJĂèęĄæşĺæĎŕåIJřæŸŕïijŇçŢĺæĹůéIJĂèęĄä£ÌèŕĄmini-batchåď ğåřŕæŸŕmicro-batchåď ğåřŔçŽĎæŢť æŢřåĂ■ãĂĆ

```
fleet.init(is_collective=True)
dist_strategy = paddle.distributed.fleet.DistributedStrategy()
strategy.pipeline_configs = {"micro_batch": 4}
```

å§žäžŐResNet50ç¡ŚçżIJçŽĎæţĄæřť çž£åźűèąŇäżčçăĄïijŽexample/resnetãĂĆ

ä¡£çŢĺäÿŃè£řåŚ¡äžď èąŇè£ŘèąŇçď žä¿ŃäżčçăĄïijŽ

```
python -m paddle.distributed.launch \
       --gpus="0,1,2,3,4" \
       train_fleet_pipeline.py
```

æŐğåĹűåŔřè¿ŞåĞžä£qæĄŕåęĆäÿŃïijŽ

```
WARNING 2021-01-08 15:53:27,677 launch.py:314] Not found distinct␣
↪arguments and compiled with cuda. Default use collective mode
launch train in GPU mode
INFO 2021-01-08 15:53:27,679 launch_utils.py:471] Local start 5␣
↪processes. First process distributed environment info (Only For␣
↪Debug):
␣
↪+=======================================================================
|                        Distributed Envs                               ␣
↪Value                   |
+-----------------------------------------------------------------
↪---------------------+
|                       PADDLE_TRAINER_ID                             0␣
↪                      |
|                 PADDLE_CURRENT_ENDPOINT                   127.0.0.
↪1:52033                |
|                     PADDLE_TRAINERS_NUM                           5␣
↪                      |
```

(äÿŃéąţçżğçż■)

(çż■äÿŁéąţ)

```
 |                 PADDLE_TRAINER_ENDPOINTS  ... 0.1:12178,127.0.0.
↪1:28915,127.0.0.1:32114|
 |                     FLAGS_selected_gpus                         0␣
↪                     |
␣
↪+=============================================================================
 INFO 2021-01-08 15:53:27,679 launch_utils.py:475] details abouts␣
↪PADDLE_TRAINER_ENDPOINTS can be found in log/endpoints.log.
 grep: warning: GREP_OPTIONS is deprecated; please use an alias or␣
↪script
 server not ready, wait 3 sec to retry...
 not ready endpoints:['127.0.0.1:40388', '127.0.0.1:12178', '127.0.
↪0.1:28915', '127.0.0.1:32114']
 server not ready, wait 3 sec to retry...
 not ready endpoints:['127.0.0.1:12178']
 W0108 15:53:37.673019 103703 device_context.cc:342] Please NOTE:␣
↪device: 0, GPU Compute Capability: 7.0, Driver API Version: 11.0,␣
↪Runtime API Version: 10.1
 W0108 15:53:37.678391 103703 device_context.cc:352] device: 0,␣
↪cuDNN Version: 7.6.
```

æŮěå£Ůä£ąæĄŕä¡■äžŐlogçŻőå¡ŢäÿŃïijŇlog/workerlog.4æŮěå£ŮæŰĞäžűçŽĎåĘĚåőźåęĆäÿŃïijŽ

```
grep: warning: GREP_OPTIONS is deprecated; please use an alias or␣
 ↪script
W0108 15:52:27.723405 103188 device_context.cc:342] Please NOTE:␣
 ↪device: 4, GPU Compute Capability: 7.0, Driver API Version: 11.0,␣
 ↪Runtime API Version: 10.1
W0108 15:52:27.728278 103188 device_context.cc:352] device: 4,␣
 ↪cuDNN Version: 7.6.
I0108 15:52:32.665313 103188 gen_nccl_id_op_helper.cc:176] Server␣
 ↪listening on: 127.0.0.1:32347 successful.
W0108 15:52:36.874132 103188 operator.cc:1194] Device index is only␣
 ↪supported under pipeline parallelism, so it will be ignored.
grep: warning: GREP_OPTIONS is deprecated; please use an alias or␣
 ↪script
W0108 15:53:31.393914 103723 device_context.cc:342] Please NOTE:␣
 ↪device: 4, GPU Compute Capability: 7.0, Driver API Version: 11.0,␣
 ↪Runtime API Version: 10.1
W0108 15:53:31.398906 103723 device_context.cc:352] device: 4,␣
 ↪cuDNN Version: 7.6.
I0108 15:53:34.465754 103723 gen_nccl_id_op_helper.cc:176] Server␣
 ↪listening on: 127.0.0.1:32114 successful.
W0108 15:53:40.784844 103723 operator.cc:1194] Device index is only␣
 ↪supported under pipeline parallelism, so it will be ignored.
[Epoch 0, batch 5] loss: 0.37770, acc1: 0.03125, acc5: 0.03125
[Epoch 0, batch 10] loss: 0.06200, acc1: 0.00000, acc5: 0.03125
[Epoch 0, batch 15] loss: 0.26105, acc1: 0.00000, acc5: 0.00000
[Epoch 0, batch 20] loss: 0.00000, acc1: 0.00000, acc5: 0.00000
```

(äÿŃéąţçżğçż■)

(çż■äÿŁéąţ)

```
[Epoch 0, batch 25] loss: 0.37330, acc1: 0.00000, acc5: 0.06250
[Epoch 0, batch 30] loss: 0.00000, acc1: 0.00000, acc5: 0.00000
[Epoch 0, batch 35] loss: 0.07487, acc1: 0.00000, acc5: 0.00000
[Epoch 0, batch 40] loss: 0.12932, acc1: 0.03125, acc5: 0.06250
[Epoch 0, batch 45] loss: 0.19604, acc1: 0.00000, acc5: 0.03125
[Epoch 0, batch 50] loss: 0.07977, acc1: 0.00000, acc5: 0.00000
[Epoch 0, batch 55] loss: 0.00000, acc1: 0.00000, acc5: 0.00000
[Epoch 0, batch 60] loss: 0.13464, acc1: 0.00000, acc5: 0.06250
[Epoch 0, batch 65] loss: 0.13940, acc1: 0.00000, acc5: 0.03125
[Epoch 0, batch 70] loss: 0.00000, acc1: 0.00000, acc5: 0.00000
[Epoch 0, batch 75] loss: 0.00000, acc1: 0.00000, acc5: 0.00000
```

æşÍæĎŕäžŃéąź

çŤśäžŐæţĄæřť'çž£åźűèąŇåřĘæÍąåđŃçŽĎåśĆæŤ¿ç¡őåĹřäÿ■åŘŇçŽĎèőąçőŮèő¿åď'ĞïijŇåŻäæ■ď'åIJÍfet

6.6 äžŇæňąåijĂåŕŚ

- TBA

6.7 æŢť'ä¡Şçď'žä¿Ń

- TBA

CHAPTER 7

ParameterServerèő■çżČ

7.1 å£ńéĂ§åijĂåğŃ

åIJÍåď ğæŢřæ■őæţłæ¡őçŽĎæŐÍåŁÍäÿŃïijŇæIJL’æăĞç■¿èő■çżČæŢřæ■őçŽĎèğĎæÍąåŔŰå¿ŮäžĘéčđéĂ§ç
Speech 2çşżçż§ä¡£çŤÍäžĘ11940åřŔæŮűçŽĎèŕ■é§şæŢřæ■őäżěåŔŁèűĚè£ĞĞ200äÿĞåŔěèąÍè£řæÌěèő■çżČèŕ■é§

äÿžäžĘæŔŘéńŸæÍąåđŃçŽĎèő■çżČæŢĹçŐĞïijŇåĹĘåÿČåijŔèő■çżČåžŤè£ŘèĂŇçŤ§ïijŇåĚűäÿ■å§žäžŐ

- **èő■çżČèŁĆçĆź**ïijŽèŕěèŁĆçĆźèť §èt’čåőŇæĹŘæŢřæ■őèŕżåŔŰãĂĄåL’■åŘŚèőąçőŮãĂĄåŔ■åŘŚæćŕåžę
- **æIJ■åŁąèŁĆçĆź**ïijŽåIJÍæŤűåĹřæL’ĂæIJL’èő■çżČèŁĆçĆźäijăæÌěçŽĎæćŕåžęåŘŐïijŇèŕěèŁĆçĆźäijŽåì

æăźæ■őåŔĆæŢřæŻt’æŰřçŽĎæŰźåijŔäÿ■åŘŇïijŇåŔŕäžěåĹĘäÿžåŘŇæ■ě/åijĆæ■ě/GeoåijĆæ■ěäÿL’çğ■ï

- **åŘŇæ■ěèő■çżČ**ïijŽæL’ĂæIJL’WorkerçŽĎè£Žåžęä£ÌæŇĄäÿĂèĞt’ïijŇå■şæŕŔèő■çżČåőŇäÿĂäÿłBatchå
- **åijĆæ■ěèő■çżČ**ïijŽäÿŐåŘŇæ■ěèő■çżČäÿ■åŘŇïijŇåIJÍåijĆæ■ěèő■çżČäÿ■äżżä¡Ţäÿd’äÿłWorkeräźŃéŰ
- **GEOåijĆæ■ěèő■çżČ**ïijŽGEOåijĆæ■ěèő■çżČæŸŕéčđæąĺçŃňæIJL’çŽĎäÿĂçğ■åijĆæ■ěèő■çżČæĺąåijŔï
  èő■çżČèŁĆçĆźéČ¡äijŽåřĘæIJňåIJřçŽĎåŔĆæŢřèőąçőŮäÿĂæňąåůőåĂij(StepéŰt’éŽŤåÿęæÌěçŽĎåŔĆæ

æIJňèŁĆåřĘéĞĞçŤÍæŐÍè■ŘécĘå§§éİđåÿÿçżŔåĚÿçŽĎæÍąåđŃwide_and_deepäÿžä¿ŃïijŇäżŃçż■åęĆä¡Ţ
APIïijĹpaddle.distributed.fleetïijL’åőŇæĹŘåŔĆæŢřæIJ■åŁąåŹÍèő■çżČäżżåŁąïijŇæIJňæňąå£ńéĂ§åijĂåğŃçŽ

7.1.1 çL’ĹælJňèęĄæśĆ

åIJÍçijŰåĘŹåĹĘåÿČåijŔèő■çżČçÍŃåžŔäźŃåL’■ïijŇçŤÍæĹůéIJĂèęĄçąőä£ÍåůšçżŔåőL’èčĚpaddlepaddle-2.0.0-cpuæĹŰpaddlepaddle-2.0.0-gpuåŘŁäžěäÿŁçL’ĹæIJňçŽĎéčđæąĺåijĂæžŘæąĘæđűãĂĆ

7.1.2 æŞ■ä¡IJæŰźæşŢ

åŔĆæŢřæIJ■åŁąåŹĺèő■çżČçŽĎå§žæIJňäžčçăĄäÿżèęĄåŇĚæŃňåęĆäÿŃåĞăäÿłéČĺåĹĘïijŽ 1. åŕijåĚěåĹĘåÿČåijŔèő■çżČéIJĂèęĄçŽĎä¿ÌèţŰåŇĚãĂĆ 2. åőŽäźL'åĹĘåÿČåijŔæĺąåijŔåźűåĹÌåğŃåŇŰåĹĘåÿČåijŔèő■çżČçŐŕåćČãĂĆ 3. åŁăè¡¡æĺąåđŃåŔŁæŢřæ■őãĂĆ 4. åőŽäźL'åŔĆæŢřæŻt'æŰřç■ŰçŢěåŔŁäijŸåŇŰåŹĺãĂĆ 5. åijĂåğŃèő■çżČãĂĆ

äÿŃéÌćåřĘéĂŘäÿĂè£ŻèąŇèőšèğčãĂĆ

åŕijåĚěä¿ÌèţŰ

```
import paddle
import os
import paddle.distributed.fleet as fleet
import paddle.distributed.fleet.base.role_maker as role_maker
```

åőŽäźL'åĹĘåÿČåijŔæĺąåijŔåźűåĹÌåğŃåŇŰåĹĘåÿČåijŔèő■çżČçŐŕåćČ

éĂŽè£ĞfleetinitæŐěåŔčïijŇçŤĺæĹůåŔŕäžěåőŽäźL'èő■çżČçŽÿåĚşçŽĎçŐŕåćČïijŇæşĺæĎŔæ■c

```
# å¡ŞåL'■åŔĆæŢřæIJ■åŁąåŹĺæĺąåijŔåŔŁæŤŕæŇĄéÌŹæĂĄåŽ¿æĺąåijŔïijŇ␣
↪åŻäæ■d'èő■çżČåL'■å£Ěéą̇æŇĞåőŽ`paddle.enable_static()`
paddle.enable_static()
role = role_maker.PaddleCloudRoleMaker()
fleet.init(role)
```

åŁăè¡¡æĺąåđŃåŔŁæŢřæ■ő

```
# æĺąåđŃåőŽäźL'åŔĆèĂČexamples/wide_and_deepäÿ■model.py
from model import net
from reader import data_reader

feeds, predict, avg_cost = net()

train_reader = paddle.batch(data_reader(), batch_size=4)
reader.decorate_sample_list_generator(train_reader)
```

åőŽäźL'åŘŇæ■ěèő■çżČ Strategy åŔŁ Optimizer

åIJÍFleet APIäÿ■ïijŇçŤĺæĹůåŔŕäžěä¡£çŤĺfleetDistributedStrategyæŐěåŔčåőŽäźL'èĞłåů

åĚűäÿ■a_syncéĂL'éąźçŤĺäžŐåőŽäźL'åŔĆæŢřæIJ■åŁąåŹĺçŻÿåĚşçŽĎç■ŰçŢěïijŇå¡ŞåĚűèćńèő¿åőŽäÿ

```
# åőŽäźL'åijĆæ■ěèő■çżČ
dist_strategy = fleet.DistributedStrategy()
dist_strategy.a_sync = True

# åőŽäźL'åŘŇæ■ěèő■çżČ
dist_strategy = fleet.DistributedStrategy()
dist_strategy.a_sync = False

# åőŽäźL'GeoåijĆæ■ěèő■çżČ, GeoåijĆæ■ěçŻőåL'■åŕłæŤŕæŇĄSGDäijŸåŇŰçőŮæşŢ
dist_strategy = fleet.DistributedStrategy()
dist_strategy.a_sync = True
dist_strategy.a_sync_configs = {"k_steps": 100}

optimizer = paddle.optimizer.SGD(learning_rate=0.0001)
optimizer = fleet.distributed_optimizer(optimizer, dist_strategy)
optimizer.minimize(model.loss)
```

åijĂåğŃèő■çżČ

åőŇæĹŘæĺąåđŃåŕŁèő■çżČç■ŰçŢěäźěåŘŐïijŇæĹŚäžňářšåŕŕäźěåijĂåğŃèő■çżČæĺąåđŃäžĘãĂĆåŻääÿż

åŕźäžŐæIJ■åŁąåŹĺèŁĆçĆźïijŇéęŰåĚĹçŤĺinit_server()æŐěåŘčåŕźåĚűè£ŻèąŇåĹİåğŃåŇŰïijŇçĎ

åŘŇæăůåŕźäžŐèő■çżČèŁĆçĆźïijŇçŤĺinit_worker()æŐěåŘčè£ŻèąŇåĹİåğŃåŇŰåŘŐïijŇ
åijĂåğŃæL'ğèąŇèő■çżČäżżåŁąãĂĆè£ŘèąŇexe.run()æŐěåŘčåijĂåğŃèő■çżČïijŇåźűå¿ŮåĹřěő■çżČäÿ■æ

```
if fleet.is_server():
    fleet.init_server()
    fleet.run_server()
else:
    exe = paddle.static.Executor(paddle.CPUPlace())
    exe.run(paddle.static.default_startup_program())

    fleet.init_worker()

    for epoch_id in range(1):
        reader.start()
        try:
            while True:
                loss_val = exe.run(program=paddle.static.default_
↪main_program(),
                                   fetch_list=[avg_cost.name])
                loss_val = np.mean(loss_val)
                print("TRAIN ---> pass: {} loss: {}\n".format(epoch_
↪id,
                                                              loss_
↪val))
        except paddle.core.EOFException:
            reader.reset()
```

(äÿŃéąţçżğçż■)

(çż■äÿŁéąţ)

```
    fleet.stop_worker()
```

è£Řèą Ňèő■çżČèĎŽæIJň

åőŽäźĽ'åőŇèő■çżČèĎŽæIJňåŘŐïijŇæĹŚäżňåřšåŕŕäžěçŤĺfleetrunæŇĞäżd'è£ŘèąŇåĹĘåÿČåijŔäżżå
worker_numåĹĘåĹńäÿžæIJ■åŁąèŁĆçĆźåŠŇèő■çżČèŁĆçĆźçŽĎæŢřéĞŕãĂĆåIJĺæIJňä¿Ńäÿ■ïijŇæIJ■åŁąè

```
fleetrun --server_num=1 --worker_num=2 train.py
```

7.2 æĂğèČ¡å§žåĞĘ

èŕůèő£éŮőéčđæąĺ Perf Repo èŐůåŔŰéčđæąĺæĂğèČ¡å§žåĞĘæŢřæ■õãĂĆ

7.3 èő¿èőąææĂİæČş

7.3.1 çżijè£ř

åŔĆæŢřæIJ■åŁąåŹĺæŸŕäÿłçijŰçĺŃæąĘæđűïijŇçŤĺäžŐæŰźä¿£åĹĘåÿČåijŔåźűèąŇçĺŃåžŔçŽĎçijŰåĘŹï
åůěäÿŽçŢŇéIJĂèęĄèő■çżČåd'ğåđŃçŽĎæIJžåŹĺå■ęäźăæĺąåđŃïijŇè£ŹäžŻæĺąåđŃåŔĆæŢřå¿Ăå¿ĂèűĚåd
äÿĂèĹňåŔĆæŢřæIJ■åŁąåŹĺæđűæđĎåęĆåŻ¿ïijĹåŐ§åŻ¿èőžæŰĞåIJřåİĂïijL'ïijŽ

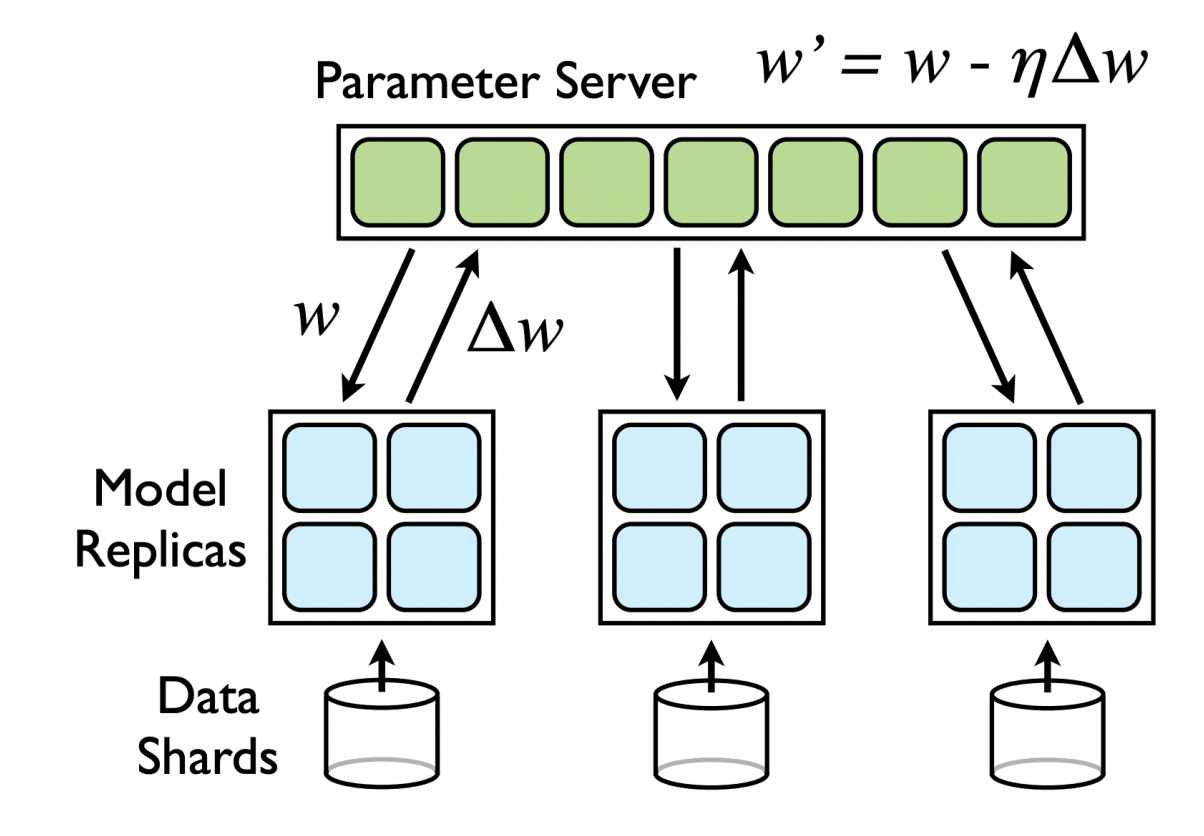

å■şåřĘæŢť'äÿłèő■çżČèŁĆçĆźåĹŠåĹĘäÿžèőąçőŮèŁĆçĆź(Worker)åŠŇåŔĆæŢřæŻť'æŰřèŁĆçĆź(PServe

éčđæąĺåŔĆæŢřæIJ■åŁąåŹĺæęĆè£ř

éčđæąĺåŔĆæŢřæIJ■åŁąåŹĺçŽĎå§žæIJňæđűæđĎæžŘèĞł Parameter Server for Distributed Machine Learning åŠŇ Large Scale Distributed Deep NetworksïijŇåźűåIJĺåĚűå§žçąĂäÿŁåĄŽäžĘåď'ğéĞŔåĹŻæŰřæİěæżąèűşçŹ¿åžęåŠŇåĚűäžŮåĚňåŔÿåŕźäžŐåŔĆæ

éčđæąĺåŔĆæŢřæIJ■åŁąåŹĺæŇěæIJL'äžěäÿŇçL'źæĂğïijŽ

1. æŤŕæŇĄåŘŇæ■ěèő■çżČãĂĄåijĆæ■ěèő■çżČåŠŇGEOåijĆæ■ěèő■çżČäÿL'çğ■æĺąåijŔïijŇèČ¡åď'§åď'Ža
2. æŤŕæŇĄå■Čäž£çžğåĹńåď'ğèğĎæĺąçĺĂçŰŔæĺąåijŔïijŇæŤŕæŇĄåĞĘåĚěãĂĄéĄŮå£Ÿç■ŰçŢěïijŇåőŇå
3. éĞĞçŤĺBRPCä¡IJäÿžåď'ŽèŁĆçĆźäźŇéŮť'éĂŽä£ąçŽĎäÿżèęĄæŰźæşŢïijŇæđĄåď'ğæŔŘå■ĞäžĘç¡ŚçżI
4. èűĚéńŸçŽĎåŘđåŘŘåŔŁåď'ŽæIJžåŁäéĂ§ærŤïijŇèČ¡åď'§æIJL'æŢĹåĹl'çŤĺèőąçőŮètĎæžŘïijŇæŔŘå■Č

éčđæąĺåŔĆæŢřæIJ■åŁąåŹĺæđűæđĎ

æđűæđĎåŻ¿åęĆäÿŇæL'Ăçď'žïijŽ

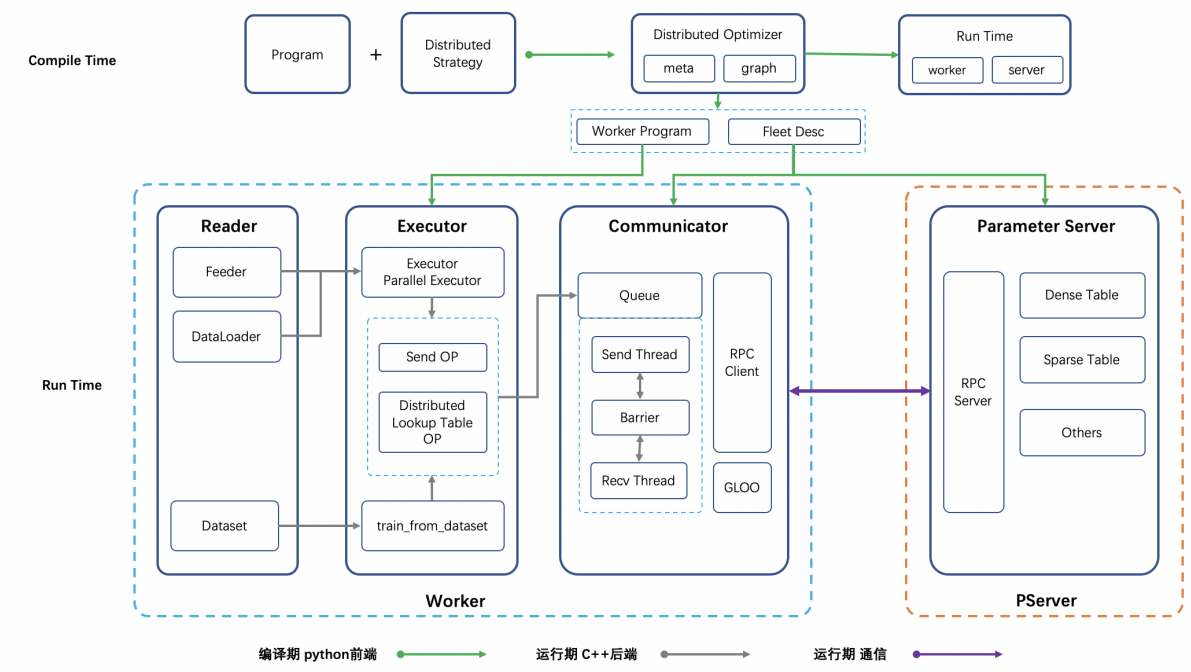

å§žæIJňçżĎäžűæŔŔè£řïijŽ

1. FleetAPI: èt'ŕçl'£æŢt'äÿłåĹęåÿČåijŔèő■çżČçŽĎAPIïijŇ åĹĘåÿČåijŔæL'ĂæIJL'åŕźåd'ŰæŽt'éIJšçŽĎAPIåİĞçŤśFleetæŽt'éIJšïijŇäÿ■åĚĄèőÿåĚűäžŰäżżä¡ŢçżĎä
2. DistributedOptimizer: çijŰèŕŚæIJ§ïijŇçżŞåŘĹéĚ■ç¡őářĘå■ŢæIJžèő■çżČç¡ŚçżIJè¡ňæ■ćäÿžåĹĘåÿČåijŔ
3. Reader: åŇĚåŘńDataset/DataLoader/FeederïijŇ Read-eräÿŐèő■çżČègčèĂęïijŇèő■çżČåŕźżäÿŐäżżæĎŔReaderéĂĆéĚ■ãĂĆ
4. Executor: æŕŔäÿłèő■çżČèŁĆçĆź(Worker)çŽĎäÿżæŰźæşŢïijŇéĂĆéĚ■åŘĎçğ■ReaderïijŇ åĹĘåÿČåijŔäÿ■åŔłéĂŽè£Ğsend/recvåŠŇåd'ŰéČĺè£ŻèąŇäžd'äžŠãĂĆ
5. Communicator: Workerçńŕèt'§èt'čæćŕåže/åŔĆæŢřèĄŽåŘĹãĂĄæŃĘåĹĘãĂĄæŤűåŔŚçŽĎæäÿå£Čæĺąåĺ
6. RPC/GLOO: èt'§èt'čåŔĆæŢřäijäéĂŠãĂĄèŁĆçĆźæŐğåĹűç■L'ïijŇéĂŽä£ąæäÿå£ČæĺąåİŮãĂĆ RPCéĂżè¿ŚäijŽäžŐæŤűåŔŚTensoræŻt'æŰřäÿżæŤűåŔŚäžŇè£ŻåĹű, GLOOèt'§èt'čæŐğåĹűèő■çżČè£ĞçÍŃäÿ■åŕźäžŐèő■çżČætĄçÍŃçŽĎæŐğåĹűïijŇåŇĚæŃňBarrierïijŇäż APIåőđçŐřåĹĘåÿČåijŔAuc/åĹĘåÿČåijŔAccç■L'éĂżè¿ŚãĂĆ
7. ParameterServer: åŔĆæŢřæIJ■åŁąåŹĺæĺąåİŮïijŇçŃňçńŇè£ŘèąŇäžŐPServerçńŕïijŇåŇĚåŘńDense/Spa

åŔĆæŢřæIJ■åŁąåŹĺæŢt'ä¡Şæđűæđ̈åĹĘçijŰèŕŚæIJ§åŠŇè£ŘèąŇæIJ§äÿd'äÿłéŸűæőţãĂĆ

çijŰèŕŚéŸűæőţïijŇæąĘæđűéIJĂåIJÍFleetAPIçŽĎéĚ■åŘĹäÿŃïijŇårĘçŤĺæĹůåőŽäźL'çŽĎå■ŢæIJžżĎç¡Ś

1. èőąçőŰèŁĆçĆź(Worker)çńŕèőąçőŮåŻ¿ïijŇ Workerçńŕç̈ŽĎèőąçőŮåŻ¿äÿżèęĄçŤšå§žçąÀèő■çżČç¡ŚçżIJ
2. éĚ■ç¡őæŰĞäżűïijŇPServerçńŕéIJÄæ■őæ■d'åŘŕåŁÍRPC ServeræIJ■åŁąïijŇäžěåŘŁçŤ§æĹŘåŘĆæŢřçŽ

è£ŘèąŇéŸűæőţïijŇPServerçńŕéIJÄåŘŕåŁÍRPCæIJ■åŁąïijŇçŻŚåŘňäźűåd'ĎçŘĘWorkerçŽĎåŘĆæŢřæŃI

åĹĘåÿČåijŔèő■çżČæĺąåijŔ

å¡ŞåL'■éčđæąĺåĚśæŤŕæŇĄäÿL'çğ■åĹĘåÿČåijŔèő■çżČæĺąåijŔïijŇåŘŇæ■ěãĂĄåijĆæ■ěãĂĄGEOåijĆæ■

- åŘŇæ■ěèő■çżČïijŽèő■çżČäÿĂäÿłminibatchåŘŐïijŇæŕŔäÿłèŁĆçĆźäijŽåŘĹåźűæĽĂæIJL'çž£çĺŃçŽĎæ PServerçńŕæŤűåĹřæĽĂæIJL'èŁĆçĆźçŽĎæćŕåžęåŘŐïijŇè£ŻèąŇæćŕåžęèĄŽåŘĹåŘŁåŘĆæŢřæŻť'æŰřã
- åijĆæ■ěèő■çżČïijŽèő■çżČäÿĂäÿłminibatchåŘŐïijŇæŕŔäÿłèŁĆçĆźçŽĎæŕŔäÿłçž£çĺŃäijŽåŘŚéĂĄæćŕã ä¡ĘæŸŕåŻääÿžåijŢåĚěäžĘåijĆæ■ěæŻť'æŰřçŽĎæIJžåĹűäijŽåŕijèĞť'èő■çżČæŢĹæđIJæIJL'æĽĂæşćåŁĺï
- GEO(Geometric Stochastic Gradient Descent)åijĆæ■ěèő■çżČïijŽGEOæŸŕéčđæąĺèĞłçăŤçŽĎåijĆæ■ěèć æŕŔäÿłèŁĆçĆźåIJĺæIJňåIJřèő■çżČèŇěáźšäÿłminibatchåŘŐ(åĚůä¡Şèő■çżČåd'ŽářŚäÿłminibatchçŤśéĚ■ æŐĺè■ŘåIJĺèń■åŘŚéĞŔãĂĄèŕ■äźĽ'åŇźéĚ■ç■Ľ'éćĘå§§è£ŻèąŇä¡£çŤĺãĂĆ

7.3.2 å■ŸåĆĺèő¿èőą

æIJňèŁĆäÿżèęĄäżŃçż■åd'ğèğĎæĺąçĺĂçŰŔåŔĆæŢřæIJ■åŁąåŹĺçŽĎå■ŸåĆĺèő¿èőąãĂĆ çěđçżŔç¡ŚçżIJèő■çżČäÿ■ïijŇåŔĆæŢřåĚśåĹĘäÿžçĺăárĘåŔĆæŢřåŠŇçĺĂçŰŔåŔĆæŢřäÿd'çğ■ïijŇ åĚűäÿ■çĺăárĘåŔĆæŢřæŇĞærŔæňąèő■çżČéČ¡éIJĂèęĄåĚĺéČĺæŻť'æŰřçŽĎåŔĆæŢřïijŇä¿ŃåęĆåĚĺè£đæŐěås fully-connectedïijĽ'çŽĎæİČéĞ■ïijĹweightïijĽ'åŠŇåĄŔç¡őïijĹbiasïijĽ'ç■Ľ'ãĂĆ çĺĂçŰŔåŔĆæŢřæŇĞærŔæňąèő■çżČäżĚéIJĂéČĺăĹĘæŻť'æŰřçŽĎåŔĆæŢřïijŇä¿ŃåęĆEmbeddingèąĺïijŇærŔ

åŐ§çŘĘ

åŔĆæŢřæIJ■åŁąåŹĺäÿ■ïijŇåŔĆæŢřçŽĎå■ŸåĆĺèő¿èőąåžŤèŕěåĹĘäÿžäÿd'éČĺåĹĘïijŇåĹĘéĚ■åŠŇå■ŸåĆ

PServeråŔĆæŢřåĹĘéĚ■

çĺăárĘåŔĆæŢřäijŽåĚĺéČĺåśŢåźşæĹŘäÿĂçżť'æŢřçżĎïijŇæŃijæŐěåIJĺäÿĂèţůåŔŸæĹŘäÿĂäÿłåd'ğçŽĎäÿ

çĺĂçŰŔåŔĆæŢřçŽĎåĹĘéĚ■æŰźåijŔäÿžåŘŰä¡ŹïijŇæŕŔäÿłidåžŤèŕěècńåĹĘéĚ■åĹřåŞłäÿłPServeräÿŁïij *id % PServer_num* èőąçőŮå¿ŮåĹřãĂĆ

PServeråŔĆæŢřå■ŸåĆĺ

æŕŔäÿłPServeräÿŁçŽĎåŔĆæŢřå■ŸåĆĺæăijåijŔåęĆäÿŃåŻ¿æĽĂçd'žïijŽ

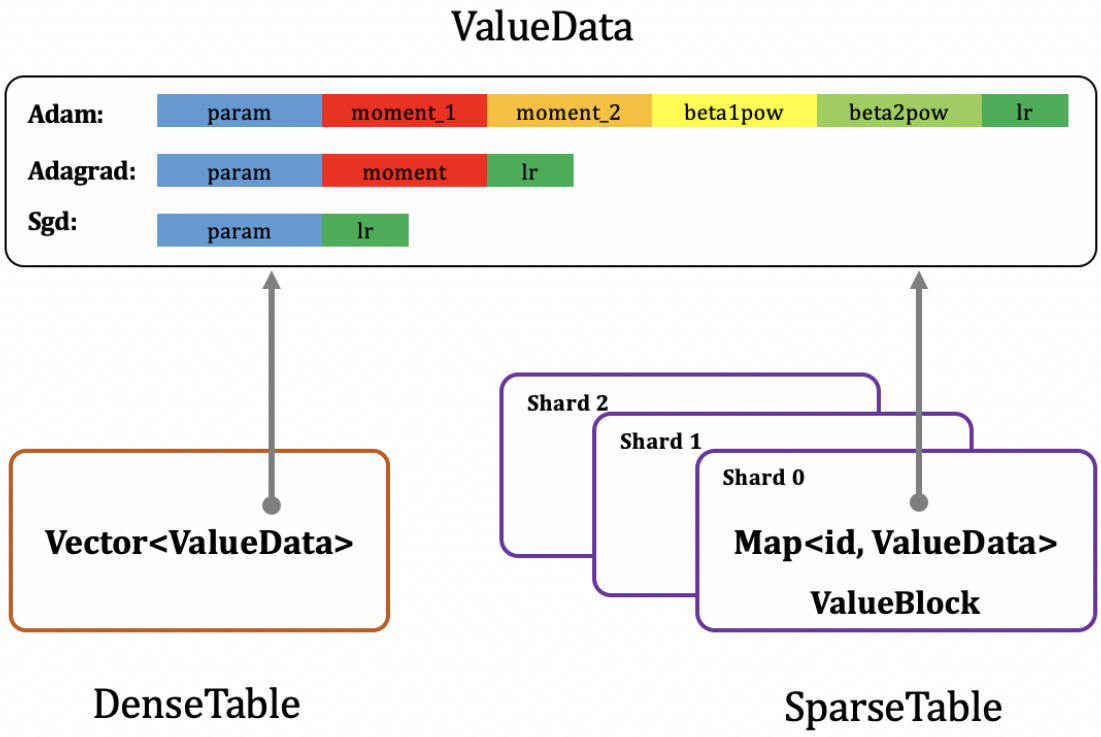

åŔŕäžěåŔŠçŐřïijŇæŮăèőžæŸŕçĺăåŕĘåŔĆæŢřïijĹDenseTableïijL'è£ŸæŸŕçĺĂçŰŔåŔĆæŢřïijĹSparseTab

èŃěäijŸåŇŰçőŮæşŢäÿžSgdïijŇéŽŔæIJžæćŕåžęäÿŃéŹ■ïijŇåŔĆæŢřæŻť'æŰřåĚňåijŔäÿžïijŽ

$$param = param - lr * grad$$

éIJĂå■ŸåĆĺåŔĆæŢřïijĹparamïijL'åŠŇå■ęäźăçŐĞïijĹlr, çżť'åžęäÿž1ïijL'ãĂĆ

èŃěäijŸåŇŰåŹĺäÿžAdagradïijŇåŔĆæŢřæŻť'æŰřåĚňåijŔäÿžïijŽ

$$moment = moment + grad * grad$$

$$param = param - \frac{lr * grad}{\sqrt{moment} + \epsilon}$$

éIJĂå■ŸåĆĺåŔĆæŢřïijĹparamïijL'ãĂĄæćŕåžęçŽĎäžŇéŸűç§l'äijřèőąïijĹmomentïijŇçżť'åžęåŠŇåŔĆæŢřäŸĂè

èŃěäijŸåŇŰåŹĺäÿžAdamïijŇåŔĆæŢřæŻť'æŰřåĚňåijŔäÿžïijŽ

$$moment\_1 = \beta_1 * moment\_1 + (1 - \beta_1) * grad$$

$$moment\_2 = \beta_2 * moment\_2 + (1 - \beta_2) * grad * grad$$

$$\beta_1{}^t = \beta_1{}^{t-1} * \beta_1$$

$$\beta_2{}^t = \beta_2{}^{t-1} * \beta_2$$

$$lr = lr * \frac{\sqrt{1 - \beta_1{}^t}}{1 - \beta_2{}^t}$$

$$param = param - lr * \frac{moment\_1}{\sqrt{moment\_2} + \epsilon}$$

éIJĂå■ŸåĆĺåŔĆæŢřïijĹparamïijL'ïijŇæćŕåžęçŽĎäÿĂéŸűãĂĄäžŇéŸűç§l'äijřèőąïijĹmoment_1,
moment_2ïijŇçżť'åžęåŠŇåŔĆæŢřäŸĂèĞť'ïijL'ïijŇäÿĂéŸűãĂĄäžŇéŸűç§l'äijřèőąçŽĎæŇĞæŢřèąřåĞŔçŐĞç
beta2pow, çżť'åžęåİĞäÿž1ïijL'äžěåŔŁå■ęäźăçŐĞïijĹlrïijŇçżť'åžęäÿž1ïijL'ãĂĆ

çĺăåŕĘåŔĆæŢřçŽĎå■ŸåĆĺæäijåijŔäÿžäÿĂäÿłäžŇçžť VectoræŢřçżĎïijŇçňňäÿĂçžť åď ğåřŔäÿžåĹĘéĚ■åĺ

äÿžäžĘèČ¡æŔŘéńŸåźűåŔŠåď ĎçŘĘèČ¡åŁŻïijŇæŕŔäÿłPServeräÿŁçĺĂçŰŔåŔĆæŢřäÿĂèĹňäijŽè£ŻèąŇå *id % shard_num* èőąçőŮå¿ŮåĹřãĂĆæŕŔäÿłshardçŽĎå■ŸåĆĺæäijåijŔäÿžå■ŮåĚÿïijĹMapïijL'ïijŇå■ŮåĚÿåĚşé

7.3.3 éĂŽä£ąèő¿èőą

7.3.4 çőĂäżŃ

çěđçżŔç¡ŚçżIJçŽĎèő■çżČè£ĞçĺŃäÿĂèĹňçŤśäÿL'éČĺåĹĘçżĎæĹŘïijŇåL'■åŘŠèőąçőŮãĂĄåŔ■åŘŠäijăæ *Loss*ïijŇçĎűåŘŐæăźæ■őéŞ¿åijŔæşŢåĹŹïijŇåŔ■åŘŠèőąçőŮå¿ŮåĹřæŕŔäÿłåŔĆæŢřçŽĎæćŕåžęïijŇæIJÅåŘŐ

åIJĺåď ŽæIJžåŔĆæŢřæIJ■åŁąåŹĺåĹĘåÿČåijŔèő■çżČäÿ■ïijŇå■ŸåIJĺäÿď çğ■äÿ■åŘŇèğŠèL'šçŽĎèŁĆçĆź åŠŇ **PServer** ïijŽ

- **Worker**ïijŽèť §èť čåőŇæĹŘæŢřæ■őèŕžåŘŰãĂĄåL'■åŘŠèőąçőŮãĂĄåŔ■åŘŠæćŕåžęèőąçőŮç■L'è£ĞçĺŇ **PServer** ãĂĆ
- **PServer**ïijŽåIJĺæŤůåĹř **Worker** äijăæİěçŽĎæćŕåžęåŘŐïijŇæăźæ■őæŇĞåőŽçŽĎäijŸåŇŰæŰźæşŢïijŇ **Worker** ïijŇåijĂåğŇæŰřäÿĂè¡őçŽĎèő■çżČãĂĆ

äÿžäžĘåĞŔåřŠéĂŽä£ąèŕůæśĆãĂĄæŔŘéńŸæŢť ä¡Şèő■çżČéĂ§åžęïijŇWorkeråIJĺårĘæćŕåžęåŘŠéĂĄçżŹ

æIJňèŁĆäÿżèęĄåŕźåĹĘåÿČåijŔèő■çżČäżżåŁąäÿ■ïijŇWorkeråŠŇPServeréŮť çŽĎéĂŽä£ąætĄçĺŃè£Żèąl

7.3.5 åŐ§çŘĘ

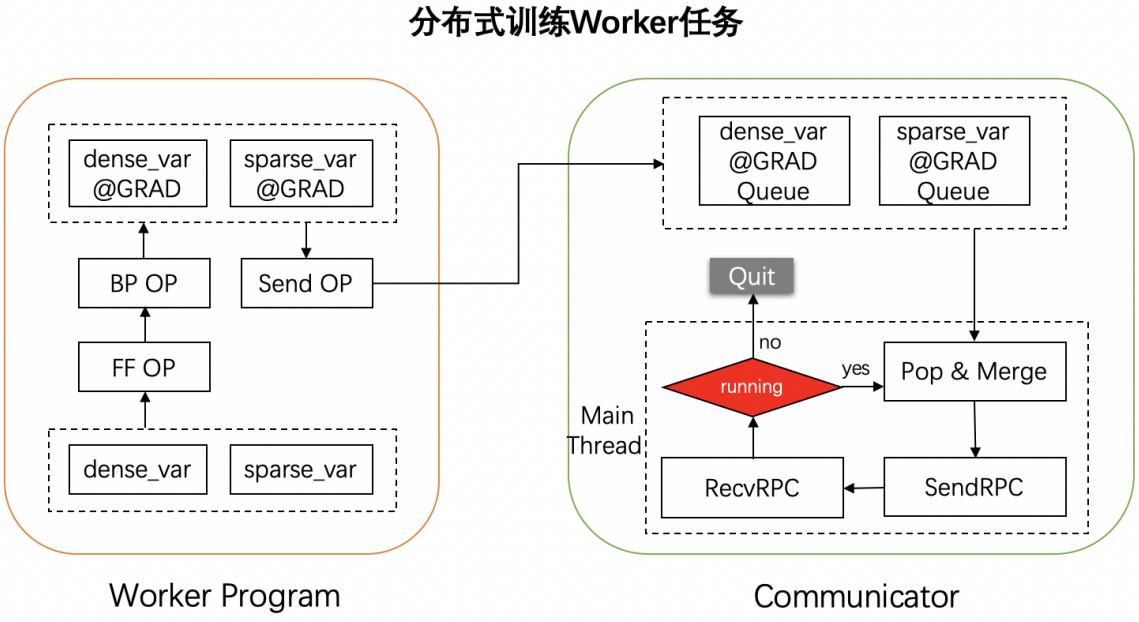

äÿŁåŻ¿åśŢçď žäžĘåĹĘåÿČåijŔèő■çżČè£ĞçĺŃäÿ■ïijŇ**Worker** äżżåŁąçŽĎæŢť äÿłætĄçĺŃãĂĆåŕŕäžěåŘŠçŐřïijŇåĹĘåÿČåijŔèő■çżČäÿ■ **Worker** äżżåŁąçŻÿè¿ČäžŐå■ŢæIJžèő■çżČäżżåŁąèĂŇèĺĂïijŇäÿżèęĄæIJL'äÿď çĆźåŇžåĹńïijŽ

1. åĹĘåÿČåijŔèő■çżČç¡ŚçżIJçŻÿè¿ČäžŐå■ŢæIJžçżĎç¡ŚèĂŇèĺĂïijŇåĹăéŹď'äžĘåŔĆæŢřæŻť'æŰřçŐůå■
   OPïijL'ãĂĆ
2. åćđåŁääžĘ **Communicator** éĂŽä£ąçżĎäžűïijŇçŤĺäżěåőŇæĹŘæćŕåžęèđ■åŘĹãĂĄåŔŚéĂĄãĂĄæŐěæŤ

CommunicatoræŸŕåĹĘåÿČåijŔåŔĆæŢřæIJ■åŁąåŹĺæąĘæđűäÿ■çŽĎéĂŽä£ąçżĎäžűïijŇçŤšåď'ŽäÿłåŔĆæ
OPïijL'åřĘæćŕåžęåŔŚéĂĄçżŹCommunicatoräÿ■æŕŔäÿłåŔĆæŢřåŕźåžŤçŽĎæćŕåžęéŸ§åĹŮéĞŇãĂĆ

CommunictoråIJĺæąĘæđűäÿ■äżěå■Ţä¿Ńå¡ćåijŔå■ŸåIJĺïijŇåIJĺ *fleet.init_worker()*
äÿ■åőŇæĹŘåĹĺåğŃåŇŰåŠŇåŘŕåŁĺãĂĆåŘŕåŁĺåŘŐïijŇCommunica-
torçŽĎäÿżçž£çĺŃäijŽäżŐæŕŔäÿłåŔĆæŢřçŽĎæćŕåžęéŸ§åĹŮäÿ■äÿ■åĄIJåŔŰåĞžæćŕåžęçŻť'èĞşæżąèűşèđ■å

- æŢřéĞŔç■L'äžŐæIJĂåď'ŽåŘŕåĚĄèőÿçŽĎæćŕåžęèđ■åŘĹæŢřéŸĹåĂijãĂĆèŕěéŸĹåĂijéĂŽè£ĞçŐŕåćČå
  FLAGS_communicator_max_merge_var_num éĚ■ç¡ő'ijŇéżŸèőď'äÿž
  CPU_NUM çŐŕåćČåŔŸéĞŔæŇĞåőŽçŽĎçž£çĺŃæŢřijŇèŃě *CPU_NUM*
  çŐŕåćČåŔŸéĞŔæIJłåőŽäźL'ïijŇåĹŹéżŸèőď'äÿž1ãĂĆ
- è£đçż■åď'Žæňąå'ĞæIJłäżŐåŔĆæŢřåŕźåžŤçŽĎæćŕåžęéŸ§åĹŮéĞŇåŔŰåĞžåĂijïijŇäÿŤåŕÌèŕŢæňąæŢřè
  FLAGS_communicator_send_wait_times éĚ■ç¡ő'ijŇéżŸèőď'äÿž5ãĂĆ

åIJĺæżąèűşæćŕåžęèđ■åŘĹæÌąäžűåŘŐïijŇåŔłèęĄå¿ÈèđåŘĹçŽĎæćŕåžęæŢřéĞŔåď'ğäžŐ0ïijŇæŮăèőžåĚ
FLAGS_communicator_max_merge_var_num ïijŇéČ¡äijŽè£ŻèąŇèđ■åŘĹãĂĄåŔŚéĂĄåŔŁæŐěæŤűãĂĆéIJĂ
distributed_lookup_table åőŇæĹŘçŽĎãĂĆ

7.4 æĂğèČ¡äijŸåŇŰ

7.4.1 èőąçőŮåŻ¿æŃĘåĹĘäÿŐäijŸåŇŰ

çőĂäżŃ

èőąçőŮåŻ¿æŃĘåĹĘçŻőåL'■äžĚæŤŕæŇĄPaddleéİŹæĂĄåŻ¿çŽĎåŔĆæŢřæIJ■åŁąåŹĺæĺąåijŔ

åŔĆæŢřæIJ■åŁąåŹĺçŽĎåĹĘåÿČåijŔèő■çżČäÿžäÿĂçğ■åÿÿèğĄçŽĎäÿ■å£ČåŇŰåĚśäžńåŔĆæŢřçŽĎåŘŇ

- **èő■çżČèŁĆçĆźïijŽèŕěèŁĆçĆźèť§èť'čåőŇæĹŘæŢřæ■őèŕžåŔŰãĂĄåL'■åŘŚèőąçőŮãĂĄåŘ■åŘŚæćŕåžę
  Worker ,äÿŐåijĆæđĎçŽĎèőąçőŮèŁĆçĆź *Heter-Worker*
- **æIJ■åŁąèŁĆçĆź**ïijŽåIJĺæŤűåĹřæL'ĂæIJL'èő■çżČèŁĆçĆźäijăæİěçŽĎæćŕåžęåŘŐïijŇèŕěèŁĆçĆźäijŽåì

  æăźæ■őåŔĆæŢřæŻť'æŰřçŽĎæŰźåijŔäÿ■åŘŇïijŇåŔŕäžěåĹĘäÿžåŘŇæ■ě/åijĆæ■ě/GeoåijĆæ■ěäÿL'çğ■ï
- **åŘŇæ■ěèő■çżČ**ïijŽåIJĺåŘŇæ■ěåŔĆæŢřæIJ■åŁąåŹĺåĹĘåÿČåijŔèő■çżČäÿ■ïijŇæL'ĂæIJL'èő■çżČèŁĆç
- **åijĆæ■ěèő■çżČ**ïijŽäÿŐåŘŇæ■ěèő■çżČäÿ■åŘŇïijŇåIJĺåijĆæ■ěèő■çżČäÿ■äżżä¡Ţäÿď'äÿłèő■çżČèŁĆçĆ
- **GEOåijĆæ■ěèő■çżČ**ïijŽäÿŐåL'■äÿď'çğ■èő■çżČäÿ■åŘŇïijŇGEOåijĆæ■ěèő■çżČäź§æŸŕäÿĂçğ■åijĆæ
  åŔ■åŘŚ-åŔĆæŢřäijŸåŇŰ)ïijŇèő■çżČè£ĞçĺŃäÿ■äżżä¡Ţäÿď'äÿłèő■çżČèŁĆçĆźäźŃéŮť'çŽĎåŔĆæŢřæŽ
  èő■çżČèŁĆçĆźäijŽåřĘæIJňåIJřçŽĎåŔĆæŢřèőąçőŮäÿĂæňąåůőåÀij(StepéŮť'éŽŤåÿęæÌěçŽĎåŔĆæŢřå

æăźæ■őèő■çżČèŁĆçĆźæL'Ăä¡£çŤĺçŽĎèő¿åď'ĞïijŇäÿŐæŢť'ä¡Şæđűæđ̃ĎçŽĎäÿ■åŘŇïijŇåŔŕäžěåĹĘäÿž
PS-CPU/PS-GPU/PS-HeteräÿL'çğ■ïijŽ

- **PS-CPU**ïijŽ*PServer* ä¡£çŤÍCPUéŻĘç¿d'æIJžåŹÍïijŇ*Worker* ä¡£çŤÍåŘŇæđĎçŽĎCPUéŻĘç¿d'æIJžåŹÍïijŇçżĎæĹŘèő■çżČæŃŞæL'Ś
- **PS-GPU**ïijŽ*PServer* ä¡£çŤÍCPUéŻĘç¿d'æIJžåŹÍïijŇ*Worker* ä¡£çŤÍåŘŇæđĎçŽĎGPUéŻĘç¿d'æIJžåŹÍïijŇçżĎæĹŘèő■çżČæŃŞæL'Ś
- **PS-Heter**ïijŽ*PServer* ä¡£çŤÍCPUéŻĘç¿d'æIJžåŹÍïijŇ*Worker* ä¡£çŤÍåŘŇæđĎçŽĎCPUéŻĘç¿d'æIJžåŹÍïijŇ*Heter-Worker* ä¡£çŤÍåijĆæđĎçŽĎAIçőŮåŁŻéŻĘç¿d'æIJžåŹÍïijĹGPU/Kunlunç■L'ïijL'ïijŇäÿL'èĂĚçżĎæĹŘèő■çżČæŃ

æIJňæŰĞåřĘåĚůä¡ŞåśŢåijĂ, èŕęçżĘäżŃçż■åŘĎäÿłèğŠèL'šçŽĎèőąçőŮåŻ¿æŃĘåĹĘåŐ§çŘĘïijŇåŔŁåęĆ

åŐ§çŘĘ

åŔĆæŢřæIJ■åŁąåŹÍçŽĎèőąçőŮåŻ¿æŃĘåĹĘïijŇæŇL'çĚğèğŠèL'šäÿ■åŘŇïijŇèőąçőŮåŻ¿äź§æIJL'æL'Ăä

æůśåžęå■ęäźăç¡ŚçżIJäÿ■æIJL'äÿd'äÿłå§žæIJňåĚČçt'ăïijŽ

- Operatorïij Žopïij Ňçż ĎæĹŘèőąçőŮçŽĎæIJĂåřŔæŞ■ä¡IJïijŇæŕŤåęĆåŁääĞŔ/FC/Embeddingæ§ěèąĺ
- Variableïij Žvarïij Ňç¡ŚçżIJå■ęäźăçŽĎåŔĆæŢřïijŇåŔĹäijŽåĹĘäÿžçĺăåŕĘåŔĆæŢřåŠŇçĺĂçŰŔåŔĆæŢřã

  – çĺăåŕĘåŔĆæŢř(Dense_Var)æŸŕæŇĞæŕŔäÿłstepéČ¡äijŽæŻt'æŰřçŽĎåŔĆæŢřïijŇæŕŤåęĆFCçŽĎV

  – çĺĂçŰŔåŔĆæŢř(Sparse_Var)æŸŕæŇĞæŕŔäÿłstepäÿ■æŸŕå£ĚéążæŻt'æŰřçŽĎåŔĆæŢřïijŇåęĆEm

å■ŢæIJžçŽĎèőąçőŮåŻ¿åęĆäÿŃæL'Ăçd'žïijŽ

- èőąçőŮåŻ¿æŃ£åĹřåŔĆæŢřçŽĎåĂij(Var)äźŃåŘŐïijŇäijŽéęŰåĚĹæL'ğèąŇåL'■åŘŚOP(FF OP)ïijŇOPåŔŕèČ¡äijŽæL'ğèąŇåIJĺäÿ■åŘŇçŽĎèő¿åd'ĞäÿŁïijŇåęĆCPU/GPU/Kunlun æĹŰåĚűäžŰAIèŁŕçL'ĞïijŇæĹŚäżňçŤĺXPUäžčæŇĞãĂĆ
- åL'■åŘŚOPèőąçőŮåőŇæĹŘåŘŐïijŇå¿ŮåĹřlossïijŇäijŽçżğçż■èőąçőŮåŔ■åŘŚOP(BP OP)ïijŇå¿ŮåĹřåŘĎäÿłåŔĆæŢřçŽĎæćŕåžę(Var_Grad)
- æŇĞåőŽSGD/Adamç■L'äijŸåŇŰåŹÍïijŇåĹl'çŤĺåŔĆæŢřçŽĎæćŕåžęïijĹVar_Gradïij L'æŻt'æŰřåŐ§ågŇå
- éĞ■åd'■äžëäÿŁæţĄçĺŃïijŇè£■äžčåŔĆæŢřçŽĎæŻt'æŰřïijŇåőđçŐřæůśåžęå■ęäźăç¡ŚçżIJçŽĎèő■çżČ

Paddle单机运行流程

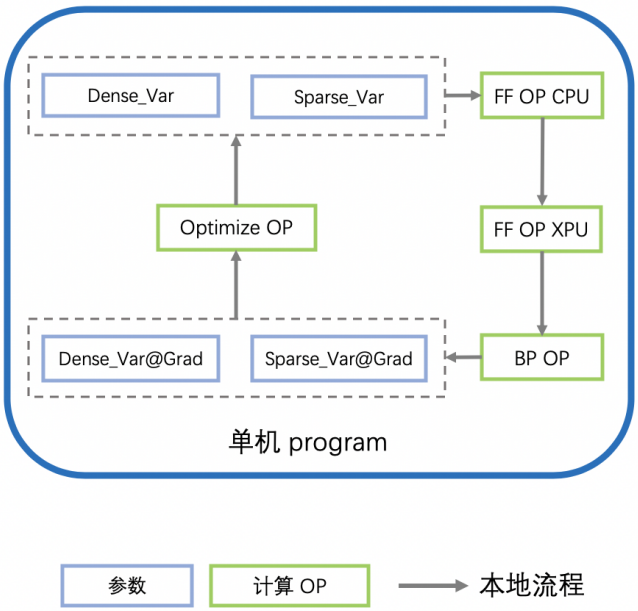

éĆčäźĹå■ŢæIJžçŽĎèőąçőŮåŻ¿åęĆä¡Ţè¡ňæ■ćäÿžåŔĆæŢřæIJ■åŁąåŹĺäÿ■åŘŇèğŠèL'šçŽĎèőąçőŮåŻ¿åŚ

èőąçőŮåŻ¿æŃĘåĹĘâĂŤâĂŤPServer

åŔĆæŢřæIJ■åŁąåŹĺæĺąåijŔäÿŃïijŇæL'ĂæIJL'åŔĆæŢřçŽĎåĚĺåśĂçIJ§åĂijéČ¡èćńåĹĘçL'Ğä£Ìå■ŸåIJĺåŘ
OPïijŇæL'ğèąŇåŔĆæŢřçŽĎæŻt'æŰřãĂĆ

äÿ■åď śäÿĂèĹňæĂğïijŇäžěPS-CPUåijĆæ■ěæĺąåijŔçŽĎèőąçőŮåŻ¿äżŃçż■PServerçŽĎèőąçőŮåŻ¿ïijŇåę

Paddle 纯CPU参数服务器Program

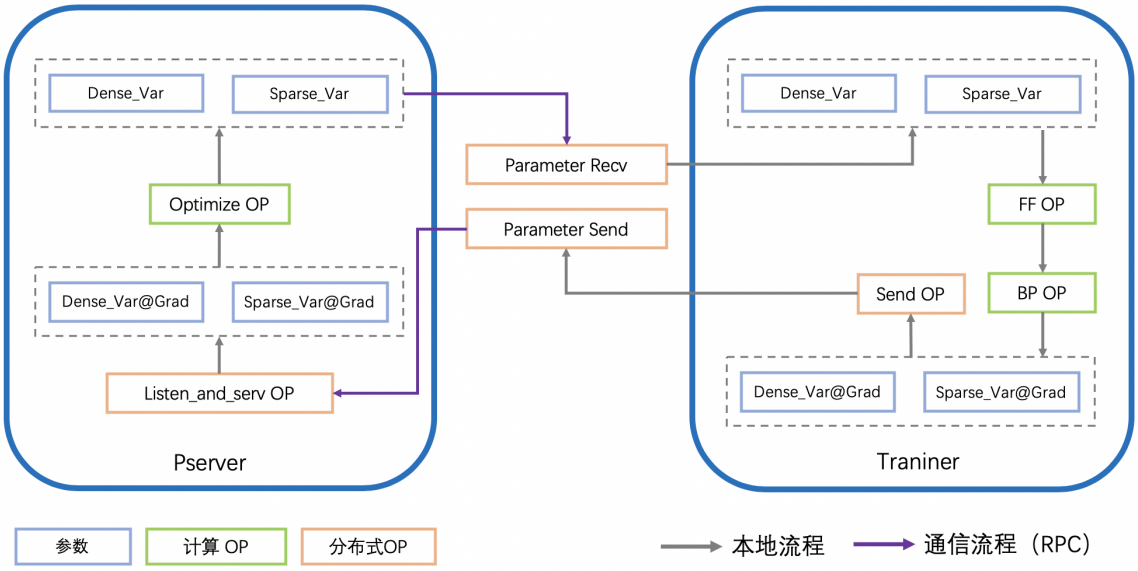

- Worker(Trainer)åIJÍèőąçŐůå¿ŮåĹřåŔĆæŢřçŽĎæćŕåžę(Var_Grad)åŘŐïijŇäijŽéĂŽè£ĞRPCåŔŠéĂĄçż
- PServerçŻŠåŘńéĂŽä£ąçńŕåŔčïijŇářĘæŤűåĹřçŽĎäÿ■åŘŇåŔĆæŢřåĹĘåĹńéĂŽè£Ğäÿ■åŘŇçŽĎOprim OPåőŇæĹŘæŻťæŰř
- WorkeråIJÍäÿŃäÿĂäÿłè£■äžčæŮűïijŇèŕůæśĆPServeräÿŁæIJĂæŰřçŽĎåŔĆæŢř
- éĞ■åď■äžëäÿŁæţĄçÍŃïijŇè£■äžčåŔĆæŢřçŽĎæŻťæŰřïijŇåőđçŐřåĹĘåÿČåijŔåŔĆæŢřæIJ■åŁąåŹÍçŽĎ

éĂŽè£ĞäÿŁè£řæţĄçÍŃïijŇPserverçŽĎèőąçŐůåŻ¿åőđçŐřçŽĎåŁ§èČ¡äÿżèęĄåĹĘäÿžäÿL'çśżïijŽ

1. æL'ğèąŇOptimizerïijŇè£ŻèąŇåŔĆæŢřæŻťæŰřçŽĎåŁ§èČ¡
2. æŐěæŤűWorkeråŘŠéĂĄçŽĎæćŕåžęïijŇèğęåŘŠOptimzerçŽĎåŁ§èČ¡
3. æŐěæŤűWorkeråŘŠéĂĄçŽĎèŕůæśĆïijŇåŘŠéĂĄæŇĞåőŽåŔĆæŢřçŽĎåŁ§èČ¡

åŁ§èČ¡2ãĂĄ3éĂŽè£ĞRPC Serverå■şåŕŕåőđçŐřïijŇæIJňèŁĆäÿ■åĘ■èţŸè£ř

åŁ§èČ¡1æIJL'äÿď'çğ■åőđçŐřéĂŤå¿ĎïijŽaãĂĄä¡£çŤÍPaddleçżĎç¡ŚïijŇæđĎæĹŘäžěoptimizer OP; bãĂĄä¡£çŤÍåőŽåĹűçŽĎæŢřæ■őçżŞæđĎåŘŁéĚ■åěŮçŽĎäijŸåŇŰçőŮæşŢåőđçŐřïijŇå■ŸåĆÍåźűæŻťæŰ

åIJÍåŘŇæ■ě/åijĆæ■ěæÍąåijŔçŽĎæČĚåĘţäÿŃïijŽ

PServerårĘèőąçŐůåŻ¿æŇL'çĚğäÿŁè£řèğĎåĹŹè£ŻèąŇçŤ§æĹŘïijŇåźűæăźæ■őèő■çżČéIJĂèęĄïijŇæůžå Decayç■L'æŞ■ä¡IJçżĎäžűãĂĆ

åIJÍGeoSGDçŽĎæČĚåĘţäÿŃïijŽ

åŔĆæŢřçŽĎæŻťæŰřOPècńæŤ¿ç¡őåIJÍWorkeräÿŁïijŇPServerèť§èť'čçż§ç■źåĚÍåśĂåŔĆæŢřïijŽæšąæIJL *ä¡ŐéćŠéĂŽä£ąåŔĆæŢřæIJ■åŁąåŹÍèő■çżČçőŮæşŢ*

äžčçăĄåőđçŐř

PServerçŽĎèőąçŐůåŻ¿çŤ§æĹŘæžŘäžčçăĄä¡■äžŐ build_pserver_program

ä¡£çŤÍFleet APIæŮűïijŇåŔĆèĂČäžěäÿŃpythonäžčçăĄïijŽserver_demo.py

```
# server_demo.py
import random
import paddle
import paddle.distributed.fleet as fleet
import paddle.distributed.fleet.base.role_maker as role_maker

paddle.enable_static()

input_data = paddle.static.data(name="sparse_input", shape=[
    None, 1], dtype="int64")
input_label = paddle.static.data(
    name="label", shape=[None, 1], dtype="int64")
label = paddle.cast(input_label, dtype="int64")

embedding = paddle.static.nn.embedding(
    input_data, is_sparse=True, size=[1000, 128])

fc1 = paddle.static.nn.fc(embedding, size=1024, activation="relu")
```

(äÿŃéąţçżğçż■)

(çż■äÿŁéąţ)

```
fc2 = paddle.static.nn.fc(fc1, size=512, activation="relu")
fc3 = paddle.static.nn.fc(fc2, size=256, activation="relu")
predict = paddle.static.nn.fc(fc3, size=2, activation="softmax")
cost = paddle.nn.functional.cross_entropy(input=predict,␣
↪label=label)

role = role_maker.PaddleCloudRoleMaker()
fleet.init(role)
strategy = fleet.DistributedStrategy()
strategy.a_sync = True
strategy.a_sync_configs = {"launch_barrier": False}

optimizer = paddle.optimizer.Adam(1e-4)
optimizer = fleet.distributed_optimizer(optimizer, strategy)
optimizer.minimize(cost)

if fleet.is_server():
    fleet.init_server()
```

```
export PSERVER_DEBUG=1
fleetrun --worker_num=1 --server_num=1 server_demo.py
cat log/serverlog.0
```

éĂŽè£ĞäžëäÿŁåŚ¡äžd'è£ŘèąŇ server_demo.py åŘŐïijŇæŮěå£ŮåžŤåŇĚåŘńäžëäÿŃçŽDè¿ŞåĞž

```
server:
 server_param {
     downpour_server_param {
     service_param {server_class: "BrpcPsServer" client_class:
↪"BrpcPsClient" service_class: "BrpcPsService" start_server_port:␣
↪0 server_thread_num: 12
     }
     downpour_table_param {table_id: 1 table_class:
↪"CommonSparseTable" shard_num: 256 type: PS_SPARSE_TABLE
     accessor {accessor_class: "CommMergeAccessor" fea_dim: 1000␣
↪embedx_dim: 128

     }
     common {name: "adam" table_name: "embedding_0.w_0" trainer_
↪num: 1 sync: false params: "Param" params: "Moment1" params:
↪"Moment2" params: "Beta1Pow" params: "Beta2Pow" params:
↪"LearningRate" dims: 128 dims: 128 dims: 128 dims: 1 dims: 1␣
↪dims: 1 initializers: "uniform_random&0&-0.0729324966669&0.
↪0729324966669" initializers: "fill_constant&0.0" initializers:
↪"fill_constant&0.0" initializers: "fill_constant&0.899999976158"␣
↪initializers: "fill_constant&0.999000012875" initializers: "fill_
↪constant&9.99999974738e-05"

     }
```

(äÿŃéąţçžğçż■)

(çż■äÿŁéąţ)

```
    }
    downpour_table_param {table_id: 0 table_class:
↪"CommonDenseTable" shard_num: 256 type: PS_DENSE_TABLE
    accessor {accessor_class: "CommMergeAccessor" fea_dim: 788738␣
↪embedx_dim: 1

    }
    common {name: "adam" table_name: "MergedDense" trainer_num: 1␣
↪sync: false params: "Param" params: "Moment1" params: "Moment2"␣
↪params: "Beta1Pow" params: "Beta2Pow" params: "LearningRate"␣
↪dims: 788738 dims: 788738 dims: 788738 dims: 1 dims: 1 dims: 1␣
↪initializers: "fill_constant&0.0" initializers: "fill_constant&0.0
↪" initializers: "fill_constant&0.0" initializers: "fill_constant&
↪0.899999976158" initializers: "fill_constant&0.999000012875"␣
↪initializers: "fill_constant&9.99999974738e-05"

    }

    }
    downpour_table_param {table_id: 2 table_class: "BarrierTable"␣
↪shard_num: 256 type: PS_OTHER_TABLE
    accessor {accessor_class: "CommMergeAccessor" fea_dim: 0␣
↪embedx_dim: 0

    }
    common {name: "" table_name: "barrier_table" trainer_num: 1␣
↪sync: false

    }

    }
    }
}
```

äžëäÿŁæŸŕserverèőąçőŮåŻ¿çŽĎéĚ■ç¡ő㤣ąæĄŕïijŇåŔŕäžëçIJŇåĹřäÿĂåĚśæIJL'3äÿłæŢřæ■őèąĺïijŽ

- 0åŔůèąĺå■ŸåĆĺäžĘDenseåŔĆæŢřïijŇçżť'åžęæŸŕçżĎç¡Śäÿ■æL'ĂæIJL'FCåśĆweightåŠŇbåŔĆæŢřçŽĎ
- 1åŔůèąĺå■ŸåĆĺäžĘSparseåŔĆæŢřïijŇä£Ĺå■ŸçŽĎåŔĆæŢræŸŕembedding_0.w_0ïijŇ çżť'åžęæŸŕ[1000, 128]ïijŇæŻť'æŰřæŰźåijŔæŸŕadam
- 2åŔůèąĺæŸŕæŐğåĹűåŘĎäÿłèŁĆçĆźéŮť'åĹİåğŃåŇŰçŽĎåŘŇæ■ě

èőąçőŮåŻ¿æŃĘåĹĘâĂŤâĂŤWorker

åŔĆæŢřæIJ■åŁąåŹĺæĺąåijŔäÿŃïijŇèő■çżČè£ĞçĺŃçŽĎæŢřæ■őèŕżåŔŰïijŇåL'■åŘŚèőąçőŮïijŇåŔ■åŘŚè

äÿ■åď śäÿĂèĹňæĂğïijŇäžěPS-CPUåijĆæ■ěæĺąåijŔçŽĎèőąçőŮåŻ¿äžŃçż■WorkerçŽĎèőąçőŮåŻ¿ïijŇåę

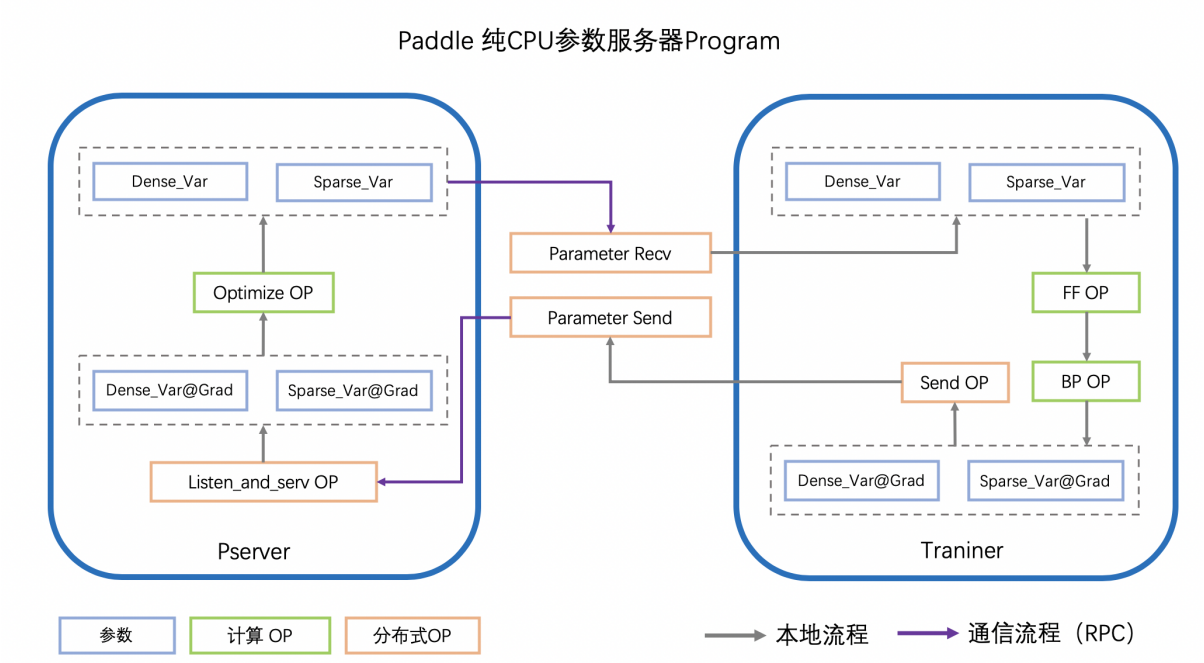

- Workerèŕżåŕ̊Ůå¡ŞåL'■æL'źæňąçŽĎèő■çżČæŢřæ■ő
- è£ŻèąŇåL'■åŘŠOPçŽĎèőąçőŮïijŇå¿ŮåĹřLoss
- å§žäžŐLossïijŇè£ŻèąŇåŕ■åŘŠOPçŽĎèőąçőŮïijŇå¿ŮåĹřåŘĎäÿłåŔĆæŢřçŽĎæćŕåžę
- åŘŠéĂĄ(Send)åŔĆæŢřçŽĎæćŕåžęçżŹPServer
- æŐěæŤű(Recv)æŻť'æŰřåŘŐçŽĎåŔĆæŢř
- éĞ■åď'■äžěäÿŁæţĄçĺŃïijŇè£■äžčèő■çżČæŢřæ■őïijŇåőđçŐřåĹĘåÿČåijŔåŔĆæŢřæIJ■åŁąåŹĺçŽĎèő■çź

  éĂŽè£ĞäÿŁè£řæţĄçĺŃïijŇWokreçŽĎèőąçőŮåŻ¿äÿŐå■ŢæIJžçŽĎèőąçőŮåŻ¿åŇžåĹńäÿžïijŽ
- åŐżéŹd'äžĘOptimzier OP
- åIJÍOptimizer OPåŐ§æİěçŽĎä¡■ç¡őåL'■ïijŇ æůżåŁääžĘSend OP
- åIJÍOptimizer OPåŐ§æİěçŽĎä¡■ç¡őåŘŐïijŇ æůżåŁääžĘRecv OP

åIJÍåŘŇæ■ě/åijĆæ■ěæląåijŔçŽĎæČĚåEţäÿŃïijŽ

WorkerçŽĎèőąçőŮåŻ¿æŇL'çĚğäÿŁè£řèğĎåĹŹè£ŻèąŇçŤ§æĹŘïijŇåźűæäźæ■őèő■çżČéIJĂèęĄïijŇæůż

çŽőåL'■PaddleçŽĎåőđçŐřäÿ■ïijŇéĂŽä£ąæţĄçĺŃä¡£çŤĺå■Ţä¿Ń *Communicator* åőđçŐřïijŇåĚĺåijĆæ■ěè£ŻèąŇèő■çżČäÿŐéĂŽä£ąïijŇåŽäæ■d'èőąçőŮåŻ¿äÿ■äžĚåIJÍæIJĂåŘŐæIJL'Send OPïijŇä¡IJçŤĺæŸŕèğęåŘŠCommunicator

åIJÍGeoSGDçŽĎæČĚåEţäÿŃïijŽ

WorkeråőđçŐřåŔĆæŢřæŻť'æŰřçŽĎåĚĺæţĄçĺŃïijŇéĂŽè£ĞSend OP èğęåŘŠ *GeoCommunicator*ïijŇèőąçőŮåźűåŘŠéĂĄæIJňåIJřäÿŐåĚĺåśĂåŔĆæŢřçŽĎdiffïijŇæŻť'åd'Žèŕęçżę̊Ęäżńçż■ïijŇåŕŕäžěå *ä¡ŐéćŠéĂŽä£ąåŔĆæŢřæIJ■åŁąåŹĺèő■çżČçőŮæşŢ*

äžčçăĄåőđçŐř

WorkerçŽĎèőąçőŮåŻ¿çŤ§æĹŘæžŘäžčçăĄä¡■äžŐ build_trainer_program

ä¡£çŤÍFleet APIæŮűïijŇåŔĆèĂČäžëäÿŃpythonäžčçăĄïijŽworker_demo.py

```
# worker_demo.py
import random
import paddle
import paddle.distributed.fleet as fleet
import paddle.distributed.fleet.base.role_maker as role_maker

paddle.enable_static()

input_data = paddle.static.data(name="sparse_input", shape=[
    None, 1], dtype="int64")
input_label = paddle.static.data(
    name="label", shape=[None, 1], dtype="int64")
label = paddle.cast(input_label, dtype="int64")

embedding = paddle.static.nn.embedding(
    input_data, is_sparse=True, size=[1000, 128])

fc1 = paddle.static.nn.fc(embedding, size=1024, activation="relu")
fc2 = paddle.static.nn.fc(fc1, size=512, activation="relu")
fc3 = paddle.static.nn.fc(fc2, size=256, activation="relu")
predict = paddle.static.nn.fc(fc3, size=2, activation="softmax")
cost = paddle.nn.functional.cross_entropy(input=predict,␣
↪label=label)

role = role_maker.PaddleCloudRoleMaker()
fleet.init(role)
strategy = fleet.DistributedStrategy()
strategy.a_sync = True
strategy.a_sync_configs = {"launch_barrier": False}

optimizer = paddle.optimizer.Adam(1e-4)
optimizer = fleet.distributed_optimizer(optimizer, strategy)
optimizer.minimize(cost)

if fleet.is_worker():
    print("worker_main_program: {}".format(
        paddle.static.default_main_program()))
```

```
fleetrun --worker_num=1 --server_num=1 worker_demo.py
cat log/workerlog.0
```

éĂŽè£ĞäžëäÿŁåŚ¡äžd'è£ŘèąŇ worker_demo.py åŘŐïijŇåŔŕäžëæL'Şå■řwork-ercŽĎåĚĺéČĺèőąçŐůåŻ¿ïijŇåŔŚçŐřåĚűäÿ■åźűæšąæIJL'AdamçŻÿåĚşOPïijŇåźűäÿŤèőąçŐůåŻ¿æIJĂåŘŐæŸí OP

```
{Out=[]} = send(inputs={X=[u'embedding_0.w_0@GRAD']}, is_sparse = 1,
↪ op_device = , op_namescope = /, op_role = 4, op_role_var = [],␣
↪send_varnames = [u'embedding_0.w_0@GRAD'], table_id = 1)
```

(äÿŃéąţçžğçż■)

(çż■äÿŁéąţ)

```
{Out=[]} = send(inputs={X=[u'fc_0.b_0@GRAD', u'fc_0.w_0@GRAD', u'fc_
↪1.b_0@GRAD', u'fc_1.w_0@GRAD', u'fc_2.b_0@GRAD', u'fc_2.w_0@GRAD',
↪ u'fc_3.b_0@GRAD', u'fc_3.w_0@GRAD']}, is_sparse = 0, op_device =␣
↪, op_namescope = /, op_role = 4, op_role_var = [], send_varnames␣
↪= [u'Dense@Grad'], table_id = 0)
```

èőąçőŮåŻ¿æŃĘåĹĘâĂŤâĂŤHeter-Worker

åijĆæđĎåŔĆæŢřæIJ■åŁąåŹĺæĺąåijŔäÿŃïijŇèő■çżČè£ĞçĺŃçŽĎåĽ■åŘŠèőąçőŮïijŇåŔ■åŘŠèőąçőŮäÿ■ç WorkeräÿŁæĽğèąŇãĂĆ

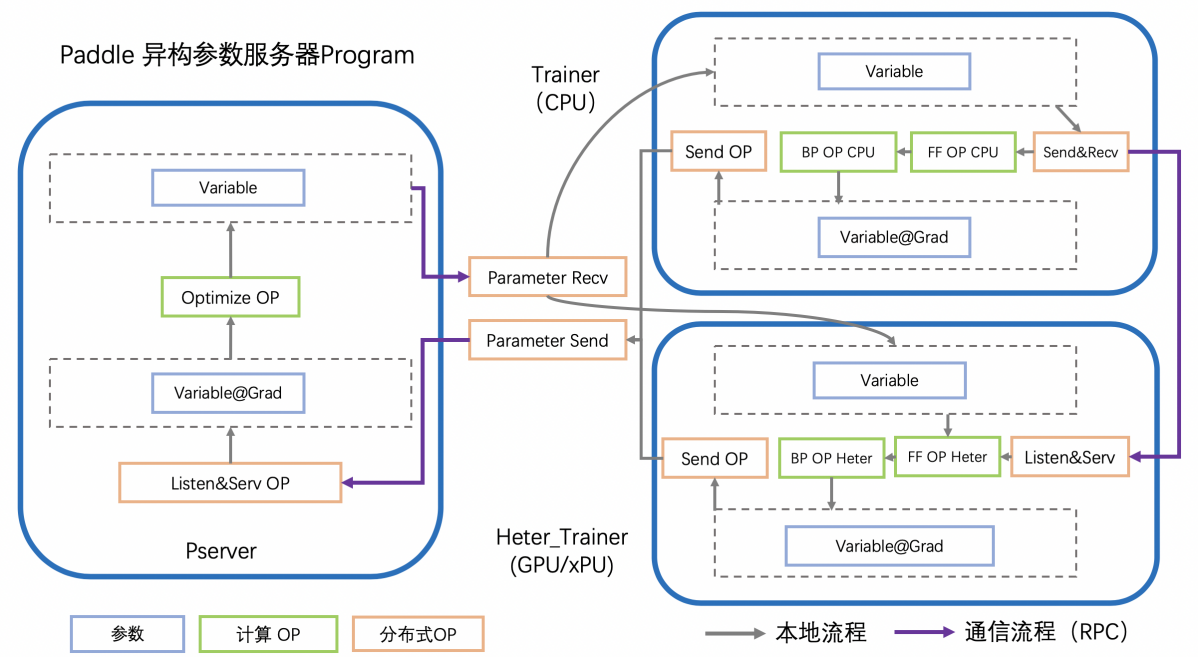

- Worker(Trainer)èŕżåŔŰå¡ŞåĽ■æĽźæňąçŽĎèő■çżČæŢřæ■ő
- WorkerèőąçőŮåĽ■ç¡őçŽĎåIJÍCPUäÿŁçŽĎåĽ■åŘŠOP
- WokerårĚåĽ■åŘŠOPèőąçőŮçżŞæđIJåŔŠçżŹHeter-Woker
- Heter-WorkerèőąçőŮåIJÍxPUäÿŁçŽĎåĽ■åŘŠOPïijŇå¿ŮåĹřloss
- Heter-WorkerèőąçőŮåIJÍxPUäÿŁçŽĎåŔ■åŘŠOPïijŇå¿ŮåĹřçŻÿåĚşåŔĆæŢřæćŕåžę
- Heter-WorkerårĚéČÍåĹĘæćŕåžęåŔŠéĂĄåŻđWorker
- WokerèőąçőŮåIJÍCPUäÿŁçŽĎåŔ■åŘŠOPïijŇå¿ŮåĹřçŻÿåĚşåŔĆæŢřæćŕåžę
- WokeräÿŐHeter-WorkeråŔŠéĂĄ(Send)æŇĄæIJL'çŽĎåŔĆæŢřçŽĎæćŕåžęçżŹPServer
- WorkeräÿŐHeter-WokeræŐěæŤű(Recv)æŻťæŰŕåŘŐçŽĎåŔĆæŢř
- éĞ■åđ'■äžëäÿŁæţĄçĺŃïijŇè£■äžčèő■çżČæŢřæ■őïijŇåőđçŐřåĹĘåÿČåijŔåŔĆæŢřæIJ■åŁąåŹĺçŽĎèő■ç éĂŽè£ĞäÿŁè£ŕæţĄçĺŃïijŇHeter-WorkeråőđçŐřçŽĎäÿżèęĄåŁ§èČ¡æŸŕïijŽ

- äÿŐWorkeréĂŽä£ąïijŇæŐěæŤűåźűåŔŠéĂĄæŇĞåőŽåŔĆæŢř
- äÿŐPServeréĂŽä£ąïijŇåŔŠéĂĄæćŕåžęïijŇæŐěæŤűæŻťæŰř
- æL'ğèąŇåL'■åŘŠ/åŔ■åŘŠ OPçŽĎè£ŘèąŇ

Heter-WokerçŽĎèőąçőŮåŻ¿çŤśäÿĂäÿłæĹŰåd'Žäÿłåij ĆæđĎBlockæđĎæĹŘïijŇæŕŔäÿłBlockäÿžäÿĂæő

äÿĂäÿłåijĆæđĎblockçŽĎè£ŘèąŇïijŇå£ĚçĎűéIJĂèęĄä¿ÌètŰåL'■ç¡őVariableçŽĎäžğåĞžïijŇåŘŇæŮűåŘ

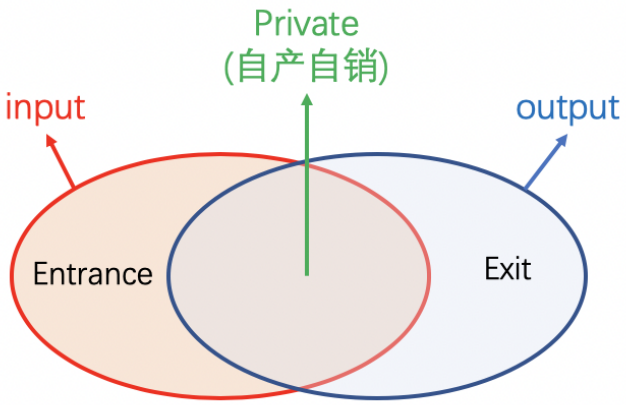

åIJĺHeter-PSæĺąåijŔäÿ■ïijŇWorkerä¡£çŤĺsend_and_recv OPæİěèğęåŔŠHetet-WorkeräÿŁçŽĎåijĆæđĎblockçŽĎè£ŘèąŇïijŇWorkeråŘŠHeter-WorkeråŘŠéĂĄEntrance VariableïijŇåŘŇæŮűç■L'å¿ĚæŐěæŤűExit Varib-aleïijŇåőđçŐřæŢť'äÿłèőąçőŮåŻ¿æţĄçĺŃçŽĎéĂŽèůřãĂĆ

Heter-PSçŽőåL'■äżĚæŤŕæŇĄAsyncæĺąåijŔ

äżčçăĄåőđçŐř

Heter-WorkerçŽĎèőąçőŮåŻ¿çŤ§æĹŘæžŘäżčçăĄä¡■äžŐ build_trainer_program

ä¡£çŤĺFleet APIæŮűïijŇåŔĆèĂČäżěäÿŃpythonäżčçăĄïijŽheter_demo.py(éIJĂèęĄåőL'èčĚGPUçL'ĹæIJ

```
# heter_demo.py
import random
import paddle
import paddle.distributed.fleet as fleet
import paddle.distributed.fleet.base.role_maker as role_maker

paddle.enable_static()

with paddle.static.device_guard("cpu"):
    input_data = paddle.static.data(name="sparse_input", shape=[
        None, 1], dtype="int64")
    input_label = paddle.static.data(
        name="label", shape=[None, 1], dtype="int64")
    label = paddle.cast(input_label, dtype="int64")
    embedding = paddle.static.nn.embedding(
        input_data, is_sparse=True, size=[1000, 128])
```

(äÿŃéąçżğçż■)

(çż■äÿŁéąţ)

```
with paddle.static.device_guard("gpu"):
    fc1 = paddle.static.nn.fc(embedding, size=1024, activation="relu
 ↪")
    fc2 = paddle.static.nn.fc(fc1, size=512, activation="relu")
    fc3 = paddle.static.nn.fc(fc2, size=256, activation="relu")
    predict = paddle.static.nn.fc(fc3, size=2, activation="softmax")
    cost = paddle.nn.functional.cross_entropy(input=predict,␣
 ↪label=label)

role = role_maker.PaddleCloudRoleMaker()
fleet.init(role)
strategy = fleet.DistributedStrategy()
strategy.a_sync = True
strategy.a_sync_configs = {"heter_worker_device_guard": "gpu",
 ↪"launch_barrier": False}

optimizer = paddle.optimizer.Adam(1e-4)
optimizer = fleet.distributed_optimizer(optimizer, strategy)
optimizer.minimize(cost)

if fleet.is_server():
    if role._is_heter_worker():
        print("heter_main_program: {}".format(
            paddle.static.default_main_program()))
```

```
fleetrun --worker_num=1 --server_num=1 --heter_worker_num=1 heter_
 ↪demo.py
cat log/heterlog.0
```

éĂŽè£ĞäžëäÿŁåŚ¡äžd'è£ŘèąŇ heter_demo.py åŘŐïijŇåŕŕäžëæL'Şå■řheter-workerçŽĎåĚĺéČĺèőąçőŮåŻ¿ïijŇåŕŕäžëåŔŚçŐřèőąçőŮåŻ¿äÿ■åŇĚåŘńäžĘäÿĂäÿłåijĆædĎblockïijŇèŕëåijĆæOPçżŞæİ§ãĂĆèŇëåŘŇæŮűæL'Şå■řworkerçŽĎèőąçőŮåŻ¿ïijŇäijŽègĆåŕ§åĹřåŐ§æİęçŽĎFCåśĆïijŇèćńsend_OPæŻ£äžčãĂĆ

äijŸåŇŰ

èğęåŔŚçĺĂçŰŕåŔĆæŢřæŻťæŰř

çĺăåŕĘåŔĆæŢřäijŽézŸèŏď'æL'ŞåŇĚäÿžäÿĂäÿłåď'ğåŔĆæŢřåŘŐïijŇåĹEçL'ĞæŤ¿åĹřåŘĎäÿłPServeräÿŁ

çĺĂçŰŕåŔĆæŢřäijŽèćńåİĞåĹĘåĹřåŘĎäÿłPServeräÿŁ

çĺĂçŰŕåŔĆæŢřçŽĎä£Ìå■ŸåŠŇæŻť'æŰřçŻőåL'■åŔŕäžëéĂŽè£ĞäžëäÿŃOPèğęåŔŚïijŇåőČäžňéČ¡åőđçŐ*distributed_lookup_table* OPïijŽ

- *paddle.nn.Embedding*

```
import paddle
paddle.enable_static()

# sparse=True,␣
↪èğęåŔŠåŔĆæŢřçŽĎçĺĂçŰŔåŇŰïijŇåŁăå£ńèő■çżČåŠŇéĂŽä£ąéĂŞåžę
embedding = paddle.nn.Embedding(
            input=x,
            size=[id_num, id_value_shape],
            sparse=True)
```

- *paddle.static.nn.embedding*

```
import paddle
paddle.enable_static()

# is_sparse=True,␣
↪èğęåŔŠåŔĆæŢřçŽĎçĺĂçŰŔåŇŰïijŇåŁăå£ńèő■çżČåŠŇéĂŽä£ąéĂŞåžę
embedding = paddle.static.nn.embedding(
            input=x,
            size=[id_num, id_value_shape],
            is_sparse=True)
```

- *paddle.fluid.layers.embedding*

```
import paddle
paddle.enable_static()

# is_sparse=True,␣
↪èğęåŔŠåŔĆæŢřçŽĎçĺĂçŰŔåŇŰïijŇåŁăå£ńèő■çżČåŠŇéĂŽä£ąéĂŞåžę
embedding = paddle.fluid.layers.embedding(
            input=x,
            size=[id_num, id_value_shape],
            is_sparse=True)
```

- *paddle.fluid.contrib.sparse_embedding*

```
import paddle
paddle.enable_static()

# sparse_embedding èğęåŔŚembçŽĎåd'ğèğĎæĺąçĺĂçŰŔ
embedding = paddle.fluid.contrib.sparse_embedding(
            input=x,
            size=[id_num, id_value_shape])
```

7.4.2 ä¡£çŤĺInMemoryDataset/QueueDatasetè£ŻèąŇèő■çżČ

æşĺæĎŔ

æIJňæŢŹçĺŃçŻőåL'■äÿ■æŤŕæŇĄåŁĺæĂĄåŻ¿ïijŃäżĚæŤŕæŇĄåIJĺpaddleéİŹæĂĄåŻ¿æĺąåijŔäÿŃä¡£çŤĺ

```
paddle.enable_static()
```

çőĂäžŃ

äÿžäžĘèČ¡énŸéĂ§è£ŘèąŇæĺąåđŃçŽĎèő■çżČïijŇæĹŚäžñä¡£çŤÍInMemoryDataset/QueueDatasetAPIè£ŻèąŇénŸæĂğèČ¡çŽĎIOïijŇåĚůä¡ŞäžŃçż■åŕŕäžěåŔĆèĂČæŰĞæąčInMemoryDatåŠŇ QueueDataset, äžěäÿŃçőĂçğřDatasetãĂĆDatasetæŸŕäÿžåď Žçž£çÍŃåŔŁåĚÍåijĆæ■ěæŰźåijŔéĞŔèžråď ŽæűĹèť źèĂĚçŽĎæĺąåijŔïijŇäijŽæđĄåď ğçŽĎåŁäéĂ§æĹŚäžñçŽĎæĺąåđŃèő■çżČãĂĆ

æIJňæŰĞäžěèő■çżČwide&deepæĺąåđŃäÿžä¿ŃïijŇåIJÍèő■çżČäÿ■åijŢåĚěå§žäžŐDataset äžěäÿŃæŸŕä¡£çŤÍDatasetæŐěåŔčäÿĂäÿłærŤè¿ČåőŇæŢť çŽĎæţĄçÍŃïijŽ

åijŢåĚědataset

1. éĂŽè£ĞdatasetÂ = paddle.distributed.InMemoryDataset() æĹŰèĂĚdataset = paddle.distributed.QueueDataset()åĹŻåžžäÿĂäÿłDatasetåŕźèśą
2. æŇĞåőŽdatasetèŕżåŔŰçŽĎèő■çżČæŰĞäžűçŽĎåĹŮèąÍïijŇéĂŽè£Ğset_filelistéĚ■ç¡őãĂĆ
3. éĂŽè£Ğdataset.init() api è£ŻèąŇDatasetçŽĎåĹİåğŃåŇŰéĚ■ç¡őïijŇinit()æŐěåŔčæŐěæŤűèŕęèğĄapiæŰĞæąčïijŇåĹŮäÿ¿åĞääÿłéĚ■ç¡őçŽĎåĹİåğŃåŇŰ
   a. åřĘæĹŚäžñåőŽäźL’åě¡çŽĎæŢřæ■őè¿ŞåĚěæäijåijŔäijăçżŹDataset, éĂŽè£Ğuse_varéĚ■ç¡őãĂĆ
   b. æŇĞåőŽæĹŚäžñçŽĎæŢřæ■őèŕżåŔŰæŰźåijŔïijŇçŤśmy_data_generator.pyåőđçŐřæŢřæ■őèŕżåŔŰçŽĎèğĎåĹŹïijŇåŘŐéİćåřĘäijŽäžŃçż■èŕżåŔŰèğĎåĹŹçŽĎåőđçŐř, éĂŽè£ĞpipecommandéĚ■ç¡őãĂĆpipecommandæŸŕDatasetçL’źæIJL’çŽĎéĂŽè£ĞçőąéĄŞ
   c. æŇĞåőŽæŢřæ■őèŕżåŔŰçŽĎbatch_sizeïijŇéĂŽè£Ğbatch_sizeéĚ■ç¡őãĂĆ
   d. æŇĞåőŽæŢřæ■őèŕżåŔŰçŽĎçž£çÍŃæŢřïijŇäÿĂèĹňèřěçž£çÍŃæŢřåŠŇèő■çżČçž£çÍŃåžŤä£İæŇĄ

```
dataset = paddle.distributed.InMemoryDataset()
batch_size = config.config["batch_size"]
thread_num = config.config["thread_num"]
dataset.init(use_var=model.inputs, pipe_command="python reader.py",␣
↪batch_size=batch_size, thread_num=thread_num)
dataset.set_filelist([config.config["train_files_path"]])
```

åęĆä¡ŢæŇĞåőŽæŢřæ■őèŕżåŔŰèğĎåĹŹ

åIJÍäÿŁæŰĞæĹŚäžñæŔŘåĽřäžĘçŤśmy_data_generator.pyåőđçŐřåĚůä¡ŞçŽĎæŢřæ■őçőąéĄŞèŕżåŔŰèğĎåĹŹïijŇéĆčäźĹïijŇæĂŐæăůäÿždatasetåĹŻåžžæŢřæ■őèŕżåäžěäÿŃæŸŕreader.pyçŽĎåĚíéČĺäžčçăĄïijŇåĚůä¡ŞæţĄçÍŃåęĆäÿŇïijŽ 1. éęŰåĚĹæĹŚäžñéIJĂèęĄåijŢåĚědata_generatorçŽĎçśžïijŇä¡■äžŐpaddle.distributed.fleet.data_generatorãĂĆ 2. åčřæŸŐäÿĂäžŻåIJÍæŢřæ■őèŕżåŔŰäÿ■äijŽçŤĺåĽřçŽĎ

3. åĹŽåžžäÿĂäÿłå■ŘçśżWideDeepDatasetReaderïijŇçžğæĽ£fleet.data_generatorçŽĎå§žçśżïijŇå§žçśżæIJĽåď'Žçğ■éĂĽæŃľïijŇåęĆæđIJæŸŕåď'Žçğ■æŢřæ■őçśżåđŃæů
4. çžğæĽ£åźűåőđçŐřå§žçśżäÿ■çŽĎgenerate_sampleåĞ¡æŢřïijŇéĂŘèąŇèŕżåŔŰæŢřæ■õãĆèŕéåĞ¡æŢ
5. åIJĺè£ŹäÿłåŔŕäžěè£■äžčçŽĎåĞ¡æŢřäÿ■ïijŇåęĆçd'žä¿ŃäžčçăĄäÿ■çŽĎdef wd_reader()ïijŇæĹŚäžňåőŽäźĽæŢřæ■õèŕżåŔŰçŽĎéĂżè¿ŚãĂĆä¿ŃåęĆåŕźäžěèąŇäÿžå■Ţä¡■çŽĎæŢřæ■ő

6. æIJĂåŘŐïijŇæĹŚäžňéIJĂèęĄåřĘæŢřæ■őæŢt'çŘĘäÿžçĽźåőŽçŽĎbatchçŽĎæăijåijŔïijŇæĽ'■èČ¡åď§èć æĹŚäžňæŮăéIJĂåĘ■åĄŽåčřæŸŐ æăźæ■õèő¿åőŽçŽĎâĂŹbatch_sizeâĂŹ, èŕéåĞ¡æŢřäijŽåIJĺgenerator_sampleåĞ¡æŢřäžğçŤ§æăůæIJňæŢřè¿¿åĽřbatch_sizeæŮűïijŇèř local_iter()ãĂĆ

7. çőĂå■ŢæÌěèŕť'ïijŇæŢřæ■õçŽĎè¿ŞåĞžéąžåžŔäÿŐæĹŚäžňåIJĺç¡ŚçżIJäÿ■åĹŽåžžçŽĎinputså£Ěéąžæ [value]),('C1',[value]),.....('label',[value])]

```
import paddle
import paddle.distributed.fleet as fleet
import os
import sys

cont_min_ = [0, -3, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0]
cont_max_ = [20, 600, 100, 50, 64000, 500, 100, 50, 500, 10, 10, 10,
 ↪ 50]
cont_diff_ = [20, 603, 100, 50, 64000, 500, 100, 50, 500, 10, 10,␣
 ↪10, 50]
hash_dim_ = 1000001
continuous_range_ = range(1, 14)
categorical_range_ = range(14, 40)


class WideDeepDatasetReader(fleet.MultiSlotDataGenerator):

    def line_process(self, line):
        features = line.rstrip('\n').split('\t')
        dense_feature = []
        sparse_feature = []
        for idx in continuous_range_:
            if features[idx] == "":
                dense_feature.append(0.0)
            else:
                dense_feature.append(
                    (float(features[idx]) - cont_min_[idx - 1]) /␣
↪cont_diff_[idx - 1])
        for idx in categorical_range_:
            sparse_feature.append(
                [hash(str(idx) + features[idx]) % hash_dim_])
        label = [int(features[0])]
        return [dense_feature]+sparse_feature+[label]

    def generate_sample(self, line):
        def wd_reader():
            input_data = self.line_process(line)
```

(äÿŃéąçžğçż■)

(çż■äÿŁéąţ)

```
            feature_name = ["dense_input"]
            for idx in categorical_range_:
                feature_name.append("C" + str(idx - 13))
            feature_name.append("label")
            yield zip(feature_name, input_data)

        return wd_reader

if __name__ == "__main__":
    my_data_generator = WideDeepDatasetReader()
    my_data_generator.set_batch(16)

    my_data_generator.run_from_stdin()
```

å£ńéĂ§èřČèŕŢDataset

æĹŚäžňåŔŕäžěèĎśçężçżĎç¡ŚæđűæđĎïijŇå■ŢçŃňéłŇèŕĄDatasetçŽĎè¿ŞåĞžæŸŕåŘęçňęåŘĹæĹŚäžňéćĎ
cat æŢřæ■őæŰĞäžű | python datasetèŕżåŔŰpythonæŰĞäžűè£Żèądatasetäžčçă
AçŽĎèřČèŕŢ

```
cat data/part-0 | python reader.py
```

è¿ŞåĞžçŽĎæŢřæ■őæăijåijŔåęĆäÿŃïijŽ 13 0.0 0.00663349917081 0.01 0.0 0.0423125 0.054 0.12 0.0 0.074 0.0 0.4 0.0 0.0 1 371155 1 846239 1 204942 1 600511 1 515218 1 906818 1 369888 1 507110 1 27346 1 698085 1 348211 1 170408 1 597913 1 255651 1 415979 1 186815 1 342789 1 994402 1 880474 1 984402 1 208306 1 26235 1 410878 1 701750 1 934391 1 552857 1 1

çŘĘæČşçŽĎè¿ŞåĞžäÿž(æĹłåŔŰäžĘäÿĂäÿłçL'Ğæőţ)ïijŽ

```
...
13 0.0 0.00663349917081 0.01 0.0 0.0423125 0.054 0.12 0.0 0.074 0.0␣
↪0.4 0.0 0.0 1 371155 1 846239 1 204942 1 600511 1 515218 1 906818␣
↪1 369888 1 507110 1 27346 1 698085 1 348211 1 170408 1 597913 1␣
↪255651 1 415979 1 186815 1 342789 1 994402 1 880474 1 984402 1␣
↪208306 1 26235 1 410878 1 701750 1 934391 1 552857 1 1
...
```

ä¡£çŤĺDatasetçŽĎäÿĂäžŻæşĺæĎŔäžŃéąź - DatasetçŽĎå§žæIJňåŐ§çŘĘïijŽåřĘæŢřæ■őprintåĽřçijŞå■Ỳ
- datasetçŻőåL'■åŔłæŤŕæŇĄåIJĺunbuntuåŔŁCentOSç■L'æăĞåĞĘLinuxçŐŕáćĂäÿŃä¡£çŤĺïijŇåIJĺw

æŢřæ■őåĞĘåď'Ğ

åőŇæŢt'æŢřæ■őäÿŃè¡¡äžěåŔŁécĎåď'ĎçŘĘäźŃåŘŐåŔŕäžěéĂL'åŔŰäÿĂäÿłpartçŽĎæŰĞäžűä¡IJäÿžden

èő■çżČ

```
import paddle
import paddle.distributed.fleet as fleet
import config
# åijĂåŘŕpaddleéİŹæĂĄåŻ¿æĺąåijŔ
paddle.enable_static()

fleet.init()

model = X.applications.Word2vec()

"""
need config loader correctly.
"""

loader = model.load_dataset_from_file(train_files_path=[config.
↪config["train_files_path"]], dict_path=config.config["dict_path"])

strategy = fleet.DistributedStrategy()
strategy.a_sync = True
optimizer = fleet.distributed_optimizer(optimizer, strategy)

optimizer.minimize(model.cost)

if fleet.is_server():
    fleet.init_server()
    fleet.run_server()

if fleet.is_worker():
    place = paddle.CPUPlace()
    exe = paddle.static.Executor(place)

    exe.run(paddle.static.default_startup_program())

    fleet.init_worker()

    distributed_training(exe, model)
    clear_metric_state(model, place)

    fleet.stop_worker()
```

åőŇæŢť'çd'žä¿ŃäžčçăĄåŘŕäžěåŘĆèĂČ FleetX/examples/wide_and_deep_dataset çŻőå¡Ţ

éĂŽè£ĞäžěäÿŁçőĂæť'ĄçŽĎäžčçăĄïijŇå■şåŘŕäžěåőđçŐřwide&deepæĺąåđŃçŽĎåď'Žçž£çĺŃåźűåŘŚèő■

7.4.3 ä¡ŐéćŚéĂŽä£ąåŘĆæŢřælJ■åŁąåŹĺèő■çżČçőŮæşŢ

çőĂäżŃ

äijŮæĽĂåŠÍç§ěïijŇåIJÍåŘŇæ■ě/åijĆæ■ěåŔĆæŢřæIJ■åŁąåŹÍåĹĘåÿČåijŔèő■çżČäÿ■WorkeræŕŔèő■çżČ
äÿžäžĘéŹ■ä¡ÕèŁĆçĆźèğĄéĂŽä£ąåŕźèő■çżČéĂ§åžęçŽĎå¡šå§■ïijŇFleetæŘŘä¿ŻäžĘäÿĂçğ■æŻť'éńŸæ

åŐ§çŘĘ

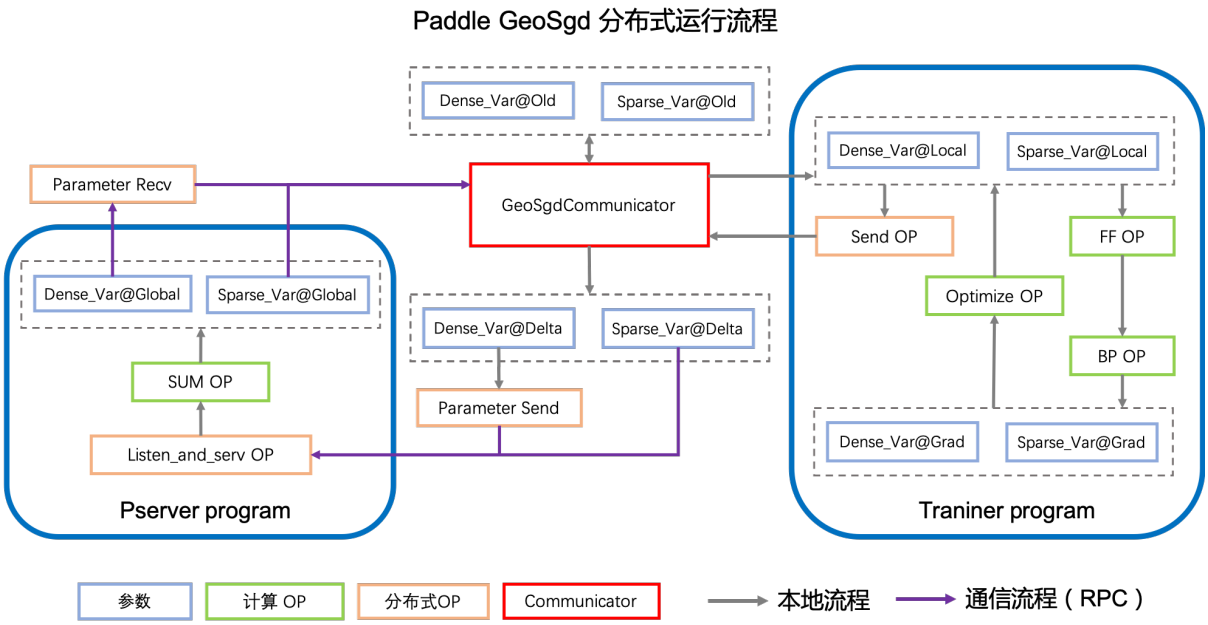

åIJÍGeoSGDæŻť'æŰřç■ŰçŢěäÿ■ïijŇWorkerçŽĎåŔĆæŢřæŻť'æŰřäź§æŸŕåIJÍåĚÍåijĆæ■ěçŽĎæİąäżűäÿŇ

- äÿŐæŹőéĂŽçŽĎåŔĆæŢřæIJ■åŁąåŹÍäÿ■åŘŇïijŇåIJÍGEOç■ŰçŢěäÿ■ïijŇæŕŔäÿłWorkerèť'§èť'čåIJÍæIJ
- GEOæŻť'æŰřç■ŰçŢěäijŽåIJÍèő■çżČè£ĞçÍŃäÿ■åŘŕåŁÍåď'Žäÿłè£ŽçÍŃïijŇèť'§èť'čåŔĆæŢřæŻť'æŰřåŔŁ

GEOç■ŰçŢěéĂŽè£ĞæÍąåđŃèő■çżČäÿŐèŁĆçĆźéĂŽä£ąåŘŇæ■ěè£ŻèąŇçŽĎæŰźåijŔïijŇåIJÍä£İèŕĄæÍą

ä¡£çŤÍæŰźæşŢ

æůżåŁää¿İèţŰ

éęŰåĚĹæĹŚäžňéIJĂèęĄæůżåŁăèő■çżČäÿ■æĽĂçŤÍåĹřçŽĎpythonæÍąåİŮpaddleåŠŇpaddle.distributed.fleetïijŇåŘŐèĂĚäÿżèęĄæŔŘä¿ŻåĹĘåÿČåijŔçŻÿåĚşçŽĎæŐěåŘčåŠŇç■ŰçŢěéĚ■ç¡őãĂ

çŻőåĽ■PaddleéżŸèőď'äÿžåŁÍæĂĄåŻ¿è£ŘèąŇæÍąåijŔïijŇåĹĘåÿČåijŔåŔĆæŢřæIJ■åŁąåŹÍèő■çżČå¡ŞåĽ

```
import paddle
import paddle.distributed.fleet as fleet
paddle.enable_static()
```

åőŽäźĽæÍąåđŃçżĎç¡Ś

åIJÍè£ŹäÿłäżŃå■Řäÿ■æĹŚäżňä¡£çŤÍWide&DeepæÍąåđŃãĂĆ

```
model = WideDeepModel()
model.net(is_train=True)
```

åĹÌåğŃåŇŰåĹĘåÿČåijŕèő■çżČçŐŕåćČ

åď ŽæIJžåŔĆæŢřæIJ■åŁąåŹÍåİĞæŸŕéĂŽè£Ğfleet.init()æŐěåŘčåĹÌåğŃåŇŰåĹĘåÿČåijŕèő■çżČ
role_maker è£ŻèąŇçŻÿåĚşéĚ■ç¡őïijŇèŃěäÿžNoneïijŇåĹŹæąĘæđűäijŽèĞłåŁÍæăźæ■őçŤÍæĹůåIJÍçŐŕåćČåŔŸ

```
fleet.init(role_maker=None)
```

éĚ■ç¡őGEOç■ŰçŢěåŔŁäijŸåŇŰçőŮæşŢ

åIJÍFleet APIäÿ■ïijŇçŤÍæĹůåŕŕäžěä¡£çŤÍfleet.DistributedStrategy()æŐěåŘčåőŽäźĽ'èĞłåů
æČşèęĄä¡£çŤÍGEOç■ŰçŢěïijŇçŤÍæĹůéęŰåĚĹéIJÄèęĄæĽŞåijĂåijĆæ■ěåŘĆæŢřæIJ■åŁąåŹÍåijĂåĚşïijI
çĎűåŘŐçŤÍæĹůéIJÄèęĄéĂŽè£ĞdistStrategy.a_sync_configsèő¿ç¡őWorkeräÿŁäijăåŔĆæŢ

```
dist_strategy = fleet.DistributedStrategy()
dist_strategy.a_sync = True
dist_strategy.a_sync_configs = {"k_steps": 400}

optimizer = paddle.optimizer.SGD(learning_rate=0.0001)

optimizer = fleet.distributed_optimizer(optimizer, dist_strategy)
optimizer.minimize(model.cost)
```

åijĂåğŃèő■çżČ

GEOç■ŰçŢěçŽĎèő■çżČäžčçăĄæš£çŤÍäžĘåŔĆæŢřæIJ■åŁąåŹÍåĹĘåÿČåijŕèő■çżČçŽĎå¡ćåijŔãĂĆ
åŕźäžŐPServerèŁĆçĆźïijŇéęŰåĚĹçŤÍinit_server()æŐěåŘčåŕźåĚűè£ŻèąŇåĹÌåğŃåŇŰïijŇçĎűåŘ
åŘŇæăůåŕźäžŐèő■çżČèŁĆçĆźïijŇçŤÍinit_worker()æŐěåŘčè£ŻèąŇåĹÌåğŃåŇŰåŘŐïijŇåijĂåğŃa

```
if fleet.is_server():
    fleet.init_server()
    fleet.run_server()
else:
    exe.run(paddle.static.default_startup_program())
    fleet.init_worker()
```

(äÿŃéąţçżğçż■)

(çż■äÿŁéąţ)

```
    # do training
    distributed_training(exe, model)
```

è£Řèą Ňæ Ű źæş Ţ

åőŇæŢť'è£ŘèąŇçd'žä¿ŃèğĄ *examples/wide_and_deep*, éIJĂæşĺæĎŘïijŇèŕěçd'žä¿ŃæŇĞåőŽçŽĎåĹĘåÿČåijŔèő■çżČæĺąåijŔäÿžåijĆæ■ěïijŇåŔŕåŔĆèĂČGEOæĺąåijŔ

éĚ■ç¡őåőŇæĹŘåŘŐïijŇéĂŽè£ĞfleetrunæŇĞäżd'è£ŘèąŇåĹĘåÿČåijŔäżžåŁąãĂĆåŚ¡äżd'çd'žä¿Ńåęworker_numåĹĘåĹńäÿžæIJ■åŁąèŁĆçĆźåŠŇèő■çżČèŁĆçĆźçŽĎæŢřéĞŔãĂĆ

```
fleetrun --server_num=2 --worker_num=2 train.py
```

7.5 åćđéĞŔèő■çżČ

7.5.1 çőĂäżŃ

åćđéĞŔèő■çżČæŸŕäÿĂçg■åÿÿèğĄçŽĎæIJžåŹĺå■ęäźăæŰźæşŢïijŇåIJĺæůšåžęå■ęäźăéćĘå§§åĘĚæIJL'åź£

éčđæąĺçŽĎåŔĆæŢřæIJ■åŁąåŹĺèő■çżČæŤŕæŇĄåćđéĞŔèő■çżČïijŇæŤŕæŇĄèő■çżČåIJĺæ§ŘäÿĂæŮűéŮ

7.5.2 åŐ§çŘĘäżŃçż■

éčđæąĺåŔĆæŢřæIJ■åŁąåŹĺåćđéĞŔèő■çżČåŇĚåŘńäÿd'éČĺåĹĘåĘĚåőźïijŇå■şæĺąåđŃä£Ìå■ŸåŠŇæĺąåđŃ

éčđæąĺæĺąåđŃåŔĆæŢřåIJĺåŔĆæŢřæIJ■åŁąåŹĺäÿŃåĹĘäÿžçĺăárĘåŔĆæŢřåŠŇçĺĂçŰŔåŔĆæŢřäÿd'çğ■ïij
åIJĺèřČçŤĺæĺąåđŃä£Ìå■ŸçŽĎæŐěåŘčåŘŐïijŇäijŽåĹĘåĹńåIJĺPServerçńŕåŠŇ0åŘůWorkerçńŕè£ŻèąŇåŔĆæŢ

éčđæąĺæĺąåđŃåŔĆæŢřåIJĺåŔĆæŢřæIJ■åŁąåŹĺäÿŃåĹĘäÿžçĺăárĘåŔĆæŢřåŠŇçĺĂçŰŔåŔĆæŢřäÿd'çğ■ïij
éIJĂèęĄåĹĘåĹńåIJĺPServerçńŕåŠŇ0åŘůWorkerçńŕè£ŻèąŇåŁăè¡¡æL'■èČ¡åőŇæĹŘåŕźåŔĆæŢřçŽĎåŁăè¡¡ãĂĆ

èő■çżČåŘŕåŁĺæŮűæŕŘäÿłPServerçŽĎå§žæIJňåĹİåğŃæţĄçĺŃåęĆäÿŃïijŽ

- æŕŘäÿłèŁĆçĆźæL'ğèąŇ *fleet.init_server(dirname=None, var_names=None, **kwargs)* è£ŻèąŇPServerçńŕåĹİåğŃåŇŰïijŇåĹĘéĚ■åĽræ■d'èŁĆçĆźçŽĎçĺăárĘåŔĆæŢřäijŽæŇL'çĚğåőŽäźL'çŽ
çĺĂçŰŔåŔĆæŢřåĹŹåŘłécĎåőŽäźL'åĞžåĹİåğŃåŇŰæŰźæşŢïijŇçĺĂçŰŔåŔĆæŢřäijŽåIJĺèő■çżČè£Ğçĺĺĺ
init_serverçŤĺæIJL'äÿd'äÿłéĂL'éĚ■åŔĆæŢřïijŇåĹĘåĹńæŸŕ
dirname'*åŠŇ*'*var_names*,'dirname'èąĺçd'žéIJĂèęĄåćđéĞŔåŁăè¡¡çŽĎæĺąåđŃèůŕå¿ĎïijŇäÿd'äÿłéĂL'éĚ
åęĆæđIJåŘłæŇĞåőŽ *dirname*ïijŇ åĹŹèąĺçd'žäijŽäźŐæŇĞåőŽçŽĎçŽő¡Ţäÿ■åŁăè¡¡åĚlélČĺçŽĎçĺĂçŰŔ
åęĆæđIJè£ŸæŇĞåőŽäžĘ'var_names'ïijŇåĹŹèąĺçd'žåŁăè¡¡æŇĞåőŽåŔĆæŢřåŘ■çŽĎçĺĂçŰŔåŔĆæŢřã
æşĺæĎŘïijŇ*init_server* åŘłäijŽåŁăè¡¡çĺĂçŰŔåŔĆæŢřïijŇçĺăárĘåŔĆæŢřçŽĎåŁăè¡¡åIJĺWorkerçńŕè£Žè

- æŕŘäÿłèŁĆçĆźæL'ğèąŇ *fleet.run_server()* èąĺæŸŐå¡ŞåL'■èŁĆçĆźåůšçżŔåĹİåğŃåŇŰæĹŘåŁ§ïijŇåŘŕä

èő■çżČåŘŕåŁĺæŮűæŕŘäÿłWorkerçŽĎå§žæIJňåĹİåğŃæţĄçĺŃåęĆäÿŃïijŽ

- æŕŕäÿłèŁĆçĆźæL'ğèąŇ *exe.run(paddle.static.default_startup_program())* è£ŻèąŇåŔĆæŢřåĹĺåğŃåŇŰãĂĆ
- 0åŔůèŁĆçĆźæL'ğèąŇ *paddle.fluid.io.load_vars()* æŇĞåőŽèęĄåŁăè¡¡çŽĎçĺăåŕĘåŔĆæŢřçŽĎåŘ■å■ŮåĹ
- æŕŕäÿłèŁĆçĆźæL'ğèąŇ *fleet.init_worker()* ïijŇ åĚűäÿ■0åŔůèŁĆçĆźçŽĎçĺăåŕĘåŔĆæŢřåřĘåŘŇæ■ěçżŹ

èĞşæ■d'ïijŇåőŇæĹŘäžĘæŢt'äÿłèő■çżČåijĂåğŃåL'■ïijŇPServeråŠŇWorkeräÿ■çĺăåŕĘåŔĆæŢřåŠŇçĺĂçU

7.5.3 åŁ§èČ¡æŢĹædlJ

- èő■çżČåijĂåğŃæŮűïijŇä¡£çŤĺäÿŁè£řæŰźæşŢïijŇåŕŕåőđçŐřæĺąåđŃåŔĆæŢřçŽĎåĚĺéĞŔåŁăè¡¡ãĂĆ
- èő■çżČçżŞæĮ§æŸŕïijŇä¡£çŤĺäÿŁè£řæŰźæşŢïijŇåŕŕåőđçŐřæĺąåđŃåŔĆæŢřçŽĎåĚĺéĞŔä£Ĺå■ŸãĂĆ

7.5.4 ä¡£çŤĺæŰźæşŢ

æĺąåđŃä£Ĺå■ŸïijŽ

```
#␣
↪åIJĺéIJÄèęĄä£Ĺå■ŸæĺąåđŃçŽĎåIJřæŰźïijŇæL'ğèąŇäÿŃéĺćçŽĎåŚ¡äżd'ïijŇå■şåŔŕåőŇæĹŘæĺ
# åĚűäÿ■ïijŇ çĺăåŕĘåŔĆæŢřäijžèćńä£Ĺå■ŸåIJĺ0åŮůWorkeräÿŁïijŇ␣
↪çĺĂçŰŔåŔĆæŢřäijžèćńä£Ĺå■ŸåIJĺærŔäÿłPServeräÿŁçŽĎåŘŇåŘ■èůŕå¿ĎäÿŃ

dirname = "/you/path/to/model"

if fleet.is_first_worker():
    fleet.save_persistables(dirname)
```

æĺąåđŃåŁăè¡¡ïijŽ

```
# æĺąåđŃåŁăè¡¡éIJÄèęĄåŇžåĹĘæŸŕPServerè£ŸæŸŕWorker
dirname = "/you/path/to/model"

if fleet.is_server():
    var_names = None
    fleet.init_server(dirname, var_names)
    fleet.run_server()

if fleet.is_worker():
    place = paddle.CPUPlace()
    exe = paddle.static.Executor(place)

    exe.run(paddle.static.default_startup_program())
    var_names = ["w", "b"]
    fluid.io.load_vars(executor=exe, dirname=path, vars=var_names)
    fleet.init_worker()
```

7.5.5 è£Řèą Ňæ Ĺ Řå Ł§æŕŘçd'ž

1. æĺąåđŃåŁăè¡¡å¡ŞåL'■åźűæšąæIJL'æŕŘçd'ž
2. æĺąåđŃä£Ìå■ŸæĹŘåŁ§ïijŇäijŽåIJÍçŻÿåžŤçŽĎçŻőå¡Ţä£Ìå■ŸäÿŃæĺąåđŃæŰĞäžűïijŇ çĺĂçŰŕåŕĆæŢřäijŽèćńä£Ìå■ŸåIJĺæŕŕäÿłPServeräÿŁçŽĎåŘŇåŘ■èůŕå¿ĎäÿŃãĂĆ

7.5.6 åÿÿèğĄéŮőécŸäÿŐæşĺæĎŕäžŃéąź

- èŁĆçĆźåŁĺæĂĄèřČæŢt'
- èő■çżČèŁĆçĆźåIJĺåŔŠçŤ§åŔŸåŇŰçŽĎæČĚåĘţäÿŃïijŇ çĺĂçŰŕåŕĆæŢřéIJĂèęĄåĄŽäÿĂæňąéĞ■æŰřåĹĘåÿČåĹĘåÿČäżěæżąèűşæŰřçŽĎåŁăè¡¡éIJĂæśĆãĂĆ
- å¡ŞåL'■æąĘæđűåźűæšąæIJL'æŔŘä¿Żæ■d'çĺĂçŰŕåŕĆæŢřéĞ■åĹĘåÿČèĎŽæIJñïijŇçŽőåL'■éIJĂèęĄçŤ
- åŁăè¡¡æŇĞåőŽçĺåårĘåŕĆæŢř
- çŤĺæĹůåŕŕäžěéĂL'æŃl'æĂğçŽĎåŁăè¡¡æL'ĂéIJĂçŽĎçĺåårĘåŕĆæŢřïijŇåĚůä¡ŞæŸŕåIJĺ 0åŔů Worker æL'ğèąŇ ‘fluid.io.load_vars‘æŮű ïijŇæŇĞåőŽçŽĎ varsçŽĎåĹÙèąĺæÌěæŐğåĹűãĂĆ
- åŁăè¡¡æŇĞåőŽçĺĂçŰŕåŕĆæŢř
- çŤĺæĹůåŕŕäžěéĂL'æŃl'æĂğçŽĎåŁăè¡¡æŇĞåőŽçŽĎçĺĂçŰŕåŕĆæŢřïijŇåĚůä¡ŞæŸŕåIJĺPServeræL'ğè

7.5.7 èőžæŰĞ/åijŢçŤĺ

[çŢě]

7.6 æţĄåijŔèő■çżČ

7.6.1 çőĂäžŃ

éčđæąĺåŕĆæŢřæIJ■åŁąåŹĺèő■çżČæŤŕæŇĄæţĄåijŔèő■çżČæĺąåijŔïijŇæŤŕæŇĄéĚ■ç¡őå■Čäž£çžğåd'ğè INT64]èŇČåŻt'åĘĚçŽĎIDæŸăărĎïijŇæŤŕæŇĄæĺąåđŃèĞłåćđéŢ£åŔŁéĚ■ç¡őçL'źå¿ĄåĞĘåĚěïijĹäÿ■å■ŸåIJ

7.6.2 åŐ§çŘĘäžŃçż■

æţĄåijŔèő■çżČ(OnlineLearning)ïijŇ å■şèő■çżČæŢřæ■őäÿ■æŸŕäÿĂæňąæĂğæŤ¿åĚěèő■çżČçşżçż§ïijŇè æŢt'äÿłèő■çżČæIJ■åŁąäÿ■åĄIJæ■ćïijŇæŢřæ■őçżŔè£ĞécĎåd'ĎçŘĘåŘŐè£ŽåĚěèő■çżČçşżçż§åŔĆäÿŐèő■çż åČŔä£ąæĄŕæţĄãĂĄåřŔèğĘécŚãĂĄçŤţåŢĘç■L'åIJžæŹŕïijŇæŕŔåd'l'éČ¡äijŽæŰřåćđåd'ğéĞŘçŽĎæŢřæ■őïijŇ èől'æŕŔåd'l'(æŕŔäÿĂåĹż)æŰřåćđçŽĎæŢřæ■őå§žäžŐäÿŁäÿĂåd'l'(äÿŁäÿĂåĹż)çŽĎæĺąåđŃè£ŻèąŇæŰřçŽĎé

7.6.3 åŁ§èČ¡æŢĹædlJ

éĂŽè£ĞåŘĹçŘĘéĚ■ç¡őïijŇåŔråőđçŐřåd'ğèğĎæĺąæţĄåijŔèő■çżČïijŇæŔŘå■ĞæŐĺè■Řçşżç§çŽĎæĂğè

æIJňæŰĞäÿ■æűL'åŔŁåĹřçŽĎçŽÿåĚşåŁ§èČ¡åŠŇä¡£çŤĺçd'žä¿ŃïijŽ 1. ä¡£çŤĺåd'ğèğĎæĺąçĺĂçŰŔçŽĎçőŮå■Řè£ŻèąŇçżĎç¡Ś 2. éĚ■ç¡őåĞĘåĚěç■ŰçŢě 3. éĚ■ç¡őæĺąåđŃä£Ìå■ŸåŔŁáćđéĞŔä£Ìå■Ÿ

7.6.4 ä¡£çŤĺæŰźæşŢ

æţĄåijŔèő■çżČæŸŕäÿłäÿŁäÿŃæÿÿçL'ţæűL'äijŮåd'ŽçŽĎèő■çżČæŰźæşŢïijŇæIJňæŰĞåŔłèt'ť'åĞžèő■çż

```
# åĹİåğŃåŇŰåĹĘåÿČåijŔçŐŕáćČ
fleet.init()

# your real net function
model = net()

# ä¡£çŤĺåŔĆæŢřæIJ■åŁąåźĺåijĆæ■ěèő■çżČæĺąåijŔ
strategy = paddle.distributed.fleet.DistributedStrategy()
strategy.a_sync = True

# åĹĘåÿČåijŔèő■çżČåŻ¿äijŸåŇŰ
adam = paddle.fluid.optimizer.Adam(learning_rate=5e-06)
adam = fleet.distributed_optimizer(adam, strategy=strategy)
adam.minimize(model.avg_cost)

# åŘŕåŁĺPServer
if fleet.is_server():
    fleet.init_server()
    fleet.run_server()

if fleet.is_worker():
    # åĹİåğŃåŇŰWorker
    exe.run(paddle.static.default_startup_program())
    fleet.init_worker()

    while True:

        # æŇĄçż■äÿ■æŰ■çŽĎäżŐ`get_ready_training_
↪set`èŐůåŔŰåŔŕèő■çżČçŽĎäźęèőŕéŻĘåŠŇçŻÿåĚşçŽĎéĚ■ç¡ő
        # äÿŃéĺćæŸŕäÿĂäÿłæŇL'åřŔæŮűèő■çżČçŽĎä¿Ńå■Ř
        dataset, hour, day = get_ready_training_dataset()

        if dataset is None:
            break

        # ä¡£çŤĺ`dataset`äÿ■çŽĎæŢřæ■őè£ŻèąŇèő■çżČåŠŇæĺąåđŃä£Ìå■Ÿ
        exe.train_from_dataset(program=paddle.static.default_main_
↪program(),                                                    (äÿŃéąţçżğçż■)
```

(çż■äÿŁéąţ)

```
                                dataset=dataset,
                                fetch_list=[model.auc],
                                fetch_info=["avg_auc"],
                                print_period=10)

        #␣
↪0åŕůä£Ìå■ŸæÍąåđŇå■şåŕŕïijŇæŕŕåd'l'çňň0äÿłåřŕæŮűè£ŻèąŇåĚÍéĞŕä£Ìå■ŸïijŇ␣
↪åL'l'ä¡ŹæŮűéŮt'è£ŻèąŇåćđéĞŕä£Ìå■Ÿ
        if fleet.is_first_worker():
            mode = 1 if hour == 0 else 2
            fleet.save_persistables(exe, "output/epoch_{}".
↪format(day), mode)

    fleet.stop_worker()
```

7.6.5 è£ŘèąŇæĹŘåŁ§æŔŘçď ž

[çŢě]

7.6.6 åÿÿèğĄéŮőécŸäÿŐæşÍæĎŔäžŃéąź

1. èő■çżČè£ĞçÍŃäÿ■ïijŇåęCéIJĂä¡£çŤÍåĹĘåÿČåijŔæŇĞæăĞïijŇèŕůåŔĆèĂČ<åĹĘåÿČåijŔæŇĞæăĞçńă
2. åęĆæđIJèő■çżČäÿ■éĂŤäÿ■æŰ■ïijŇéIJĂèęĄåŁãè¡¡æÍąåđŃåŘŐçżğçż■èő■çżČïijŇèŕůåŔĆèĂČ<åćđéĞŔ

7.6.7 èőžæŰĞ/åijŢçŤÍ

[çŢě]

7.7 åĹĘåÿČåijŔæŇĞæăĞ

7.7.1 çőĂäżŃ

åĹĘåÿČåijŔæŇĞæăĞæŸŕæŇĞåIJÍåĹĘåÿČåijŔèő■çżČäżżåŁąäÿ■çŤÍäžěèŕĐæţŃæÍąåđŃæŢĹæđIJçŽĐæŇ

7.7.2 åŐ§çŘĘ

åĹĘåÿČåijŔæŇĞæăĞçŽĐèőąçőŮäÿĂèĹňåŇĚåŘńäÿL'æ■ëïijŇäÿŃéİćæĹŚäžňäżěåĹĘåÿČåijŔåĞĘçąőçŐ

1. åĹİåğŃåŇŰåĹĘåÿČåijŔèő■çżČçŐŕåćČ

```
import paddle.distributed.fleet as fleet
fleet.init()
```

2. åőŽäźL'æŇĞæăĞèőąçőŮéIJĂèęĄçŽĎæL'ĂæIJL'äÿ■éŮt'çŁűæĂĄçż§èőąåĂijïijŇæŕŔäÿłèő■çżČèŁĆçĆ

```
...
pred, label = model()

# 1.␣
↪åőŽäźL'äÿ■éŮt'çŁűæĂĄçż§èőąåĂijïijŇæåůæIJňæĂżæŢřåŠŇæ■čçąőåĹĘçśżçŽĎæåůæI
correct_cnt = paddle.static.create_global_var(name=
↪"right_cnt", persistable=True, dtype='float32',␣
↪shape=[1], value=0)
total_cnt = paddle.static.create_global_var(name="total_
↪cnt", persistable=True, dtype='float32', shape=[1],␣
↪value=0)

# 2. èő■çżČèŁĆçĆźèĞłåůśçŽĎçŁűæĂĄçż§èőą
batch_cnt = paddle.sum(
    paddle.full(shape=[paddle.shape(label)[0], 1], fill_
↪value=1.0))
batch_accuracy = paddle.static.accuracy(input=pred,␣
↪label=label)
batch_correct = batch_cnt * batch_accuracy

paddle.assign(correct_cnt + batch_correct, correct_cnt)
paddle.assign(total_cnt + batch_cnt, total_cnt)
accuracy = correct_cnt / total_cnt
```

3. æL'ĂæIJL'èő■çżČèŁĆçĆźŮt'è£ŻèąŇ *all_reduce* æŞ■ä¡IJïijŇèŐůåŔŰåĚĺåśÄçż§èőąåĂijïijŇçĎűåŘŐæ

```
global_cnt = fleet.metrics.sum(total_cnt)
global_correct = fleet.metrics.sum(corrent_cnt)
global_accuracy = float(global_correct) / float(global_
↪cnt)
```

7.7.3 åĹĘåÿČåijŔæŇĞæăĞ

äÿžæŰźä¿£ä¡£çŤĺïijŇPaddleåIJĺ *paddle.distributed.metrics* äÿŃåřĘåÿÿèğĄçŽĎäÿĂäžŻæŇĞæăĞèőąçőŮè£ŻèąŇäžĘåřĄèčĚïijŇäÿŃéĺćåŕźè£ŹäžŻAPIçŽĎåŁ§èČ¡åŔŁåŔ

åĹĘåÿČåijŔAUC

paddle.distributed.fleet.metrics.**auc**(*stat_pos*, *stat_neg*, *scope=None*, *util=None*)

åĹĘåÿČåijŔAUCïijĹArea Under the CurveïijL'ãĂĆAUC æŸŕäÿĂäÿłäžŇåĹĘçśżäżżåŁąäÿ■åÿÿçŤĺçŽĎæŢĹæđIJèŕĎäżůæŇĞæăĞïijŇæŇĞROCæŽšçž£åŠŇæĺłåİŘæăĞè

åŔĆæŢřïijŽ

- stat_pos, (numpy.array|Tensor|string, required): å■ŢæIJžæ■čæăůä¿Ńäÿ■éŮt'çż§èőąçżŞæđIJïijŇå■şå■Ţa *paddle.static.auc* çŽĎ *stat_pos* è¿ŞåĞžãĂĆ
- stat_neg, (numpy.array|Tensor|string, required): å■ŢæIJžèt'§æăůä¿Ńäÿ■éŮt'çż§èőąçżŞæđIJïijŇå■şå■Ţa *paddle.static.auc* çŽĎ *stat_neg* è¿ŞåĞžãĂĆ
- scope, (Scope, optional)ïijŇä¡IJçŤĺå§§ïijŇèŃěäÿžNoneïijŇåĹŹä¡£çŤĺåĚĺåśĂ/éżŸèőd'ä¡IJçŤĺå§§ïijŇéż
- util, (UtilBase, optinal)ïijŇåĹĘåÿČåijŔèő■çżČåůěåĚůçśżïijŇèŃěäÿžNoneïijŇåĹŹä¡£çŤĺéżŸèőd'åůěåĚ *fleet.util*ïijŇ éżŸèőd'äÿžNoneãĂĆ

çŤĺæşŢçd'žä¿ŃïijŽ

```
...
pred, label = model()

# 1.␣
↪å■ŢæIJžçżĎç¡ŚéŸűæőţïijŇèőąçőŮæ■čèt'Śæăůä¿Ńäÿ■éŮt'çżŚèőąçżŞæđIJãĂĆ
auc, batch_auc, [batch_stat_pos, batch_stat_neg, stat_pos,␣
↪stat_neg] = \
    paddle.static.auc(input=pred, label=label)

# 2. åĹĘåÿČåijŔèő■çżČéŸűæőţïijŇèőąçőŮåĚĺåśĂAUCãĂĆ
global_auc = fleet.metrics.auc(stat_pos, stat_neg)
```

åĹĘåÿČåijŔAccuracy

paddle.distributed.fleet.metrics.**acc**(*correct*, *total*, *scope=None*, *util=None*)

åĹĘåÿČåijŔåĞEçąőçŐĞãĂĆåĞEçąőçŐĞïijĹAccuracyïijL'æŸŕåĹĘçśżäżżåŁąäÿ■åÿÿçŤĺçŽĎäÿĂäÿłæŢĹ

$$accuracy = \frac{correct}{total}$$

åĚűäÿ■ïijŇ*correct* æŸŕécĎætŃæăĞç■¿ç■L'äžŐçIJ§åőđæăĞç■¿çŽĎæăůæIJňæĂżæŢřïijŇ*total* æŸŕåĚĺéČĺæăůæIJňæĂżæŢřãĂĆ

åŔĆæŢřïijŽ

- correct, (numpy.array|Tensor|string, required): å■ŢæIJžécĎætŃæăĞç■¿ç■L'äžŐçIJ§åőđæăĞç■¿çŽĎæă
- total, (numpy.array|Tensor|string, required): å■ŢæIJžæăůæIJňæĂżæŢřãĂĆ
- scope, (Scope, optional)ïijŇä¡IJçŤĺå§§ïijŇèŃěäÿžNoneïijŇåĹŹä¡£çŤĺåĚĺåśĂ/éżŸèőd'ä¡IJçŤĺå§§ïijŇéż
- util, (UtilBase, optinal)ïijŇåĹĘåÿČåijŔèő■çżČåůěåĚůçśżïijŇèŃěäÿžNoneïijŇåĹŹä¡£çŤĺéżŸèőd'åůěåĚ *fleet.util*ïijŇ éżŸèőd'äÿžNoneãĂĆ

çŤĺæşŢçd'žä¿ŃïijŽ

```
...
pred, label = model()

# 1.␣
↪å■ŢæIJžçżĎç¡ŚéŸűæőţïijŇèőąçőŮæăůæIJňæĂżæŢřåŠŇéćĎæţŃæ■čçąőçŽĎæăůæIJňæŢř
correct_cnt = paddle.static.create_global_var(name="right_
↪cnt", persistable=True, dtype='float32', shape=[1],␣
↪value=0)
total_cnt = paddle.static.create_global_var(name="total_cnt
↪", persistable=True, dtype='float32', shape=[1], value=0)

batch_cnt = paddle.sum(
    paddle.full(shape=[paddle.shape(label)[0], 1], fill_
↪value=1.0))
batch_accuracy = paddle.static.accuracy(input=pred,␣
↪label=label)
batch_correct = batch_cnt * batch_accuracy

paddle.assign(correct_cnt + batch_correct, correct_cnt)
paddle.assign(total_cnt + batch_cnt, total_cnt)
accuracy = correct_cnt / total_cnt

# 2. åĹĘåÿČåijŔèő■çżČéŸűæőţïijŇèőąçőŮåĚĺåśĂåĞĘçąőçŐĞãĂĆ
global_accuracy = fleet.metrics.acc(correct_cnt, total_cnt)
```

åĹĘåÿČåijŔMAE

paddle.distributed.fleet.metrics.**mae**(*abserr*, *total_ins_num*, *scope=None*, *util=None*)

åĹĘåÿČåijŔåźşåİĞçżİåŕźèŕŕåůő(Mean Absolute Error)ãĂĆåźşåİĞçżİåŕźèŕŕåůőæŸŕçżİåŕźèŕŕåůőçŽĎåźşåİ *loss* æ■§åď'śåĂijãĂĆ

$$abserr = \sum |input - label|$$

$$mae = \frac{abserr}{total\_ins\_num}$$

åĚűäÿ■ïijŇ*input* æŸŕæăůæIJňéćĎæţŃçżŞæđIJïijŇ *label* æŸŕæăůæIJňçIJ§åőđæăĞç■¿ïijŇ*abserr* äÿžçżİåŕźèŕŕåůőåŠŇïijŇ*total_ins_num* æŸŕæăůæIJňæĂżæŢřãĂĆ

åŔĆæŢřïijŽ

- abserr, (numpy.array|Tensor|string, required): å■ŢæIJžçżİåŕźèŕŕåůőåŠŇçż§èőąåĂijãĂĆ
- total_ins_num, (numpy.array|Tensor|string, required): å■ŢæIJžæăůæIJňæĂżæŢřãĂĆ
- scope, (Scope, optional)ïijŇä¡IJçŤĺå§§ïijŇèŃěäÿžNoneïijŇåĹŹä¡£çŤĺåĚĺåśĂ/éżŸèőď'ä¡IJçŤĺå§§ïijŇéż

- util, (UtilBase, optinal)ïijŇåĹĘåÿČåijŔèő■çżČåůěåĚůçśżïijŇèŃěäÿžNoneïijŇåĹŹä¡£çŤĺéżŸèőď'åůěåĚ
  *fleet.util*ïijŇ éżŸèőď'äÿžNoneãĂĆ

çŤĺæşŢçd'žä¿ŃïjŽ

```
...
pred, label = model()

# 1.␣
↪å■ŢæIJžçżĎç¡ŚéŸűæőţïijŇèőąçőŮçżİåŕźèŕŕåůőåŠŇæăůæIJňæĂżæŢř
abserr = paddle.static.create_global_var(name="abserr",␣
↪persistable=True, dtype='float32', shape=[1], value=0)
total_cnt = paddle.static.create_global_var(name="total_cnt
↪", persistable=True, dtype='float32', shape=[1], value=0)

batch_cnt = paddle.sum(
    paddle.full(shape=[paddle.shape(label)[0], 1], fill_
↪value=1.0))
batch_abserr = paddle.nn.functional.l1_loss(pred, label,␣
↪reduction='sum')

paddle.assign(abserr + batch_abserr, abserr)
paddle.assign(total_cnt + batch_cnt, total_cnt)
mae = abserr / total_cnt

# 2. åĹĘåÿČåijŔèő■çżČéŸűæőţïijŇèőąçőŮåĚĺåśĂåĞEçąőçŐĞãĂĆ
global_mae = fleet.metrics.mae(abserr, total_cnt)
```

åĹĘåÿČåijŔMSE

paddle.distributed.fleet.metrics.**mse**(*sqrerr*, *ins_num*, *scope=None*, *util=None*)

åĹĘåÿČåijŔåĪĞæŰźèŕŕåůő(Mean Squared Error)ãĂĆåİĞæŰźèŕŕåůőæŸŕèŕŕåůőåźşæŰźåŠŇçŽĎåźşåİĞåĂ *loss* æ■§åď'śåĂijãĂĆ

$$sqrerr = \sum (input - label)^2$$
$$mse = \frac{sqrerr}{total\_ins\_num}$$

åĚűäÿ■ïijŇ*input* æŸŕæăůæIJňéćĎæţŃçżŞæđIJïijŇ *label* æŸŕæăůæIJňçIJ§åődæăĞç■¿ïijŇ*sqrerr* äÿžåźşæŰźèŕŕåůőåŠŇïijŇ*total_ins_num* æŸŕæăůæIJňæĂżæŢřãĂĆ

åŔĆæŢřïijŽ

- sqrerr, (numpy.array|Tensor|string, required): å■ŢæIJžåźşæŰźèŕŕåůőåŠŇçż§èőąåĂijãĂĆ
- total_ins_num, (numpy.array|Tensor|string, required): å■ŢæIJžæăůæIJňæĂżæŢřãĂĆ

- scope, (Scope, optional)ïijŇä¡IJçŤĺå§§ïijŇèŃěäÿžNoneïijŇåĹŹä¡£çŤĺåĚĺåśĂ/éżŸèőď'ä¡IJçŤĺå§§ïijŇéż
- util, (UtilBase, optinal)ïijŇåĹĘåÿČåijŕèő■çżČåůěåĚůçśżïijŇèŃěäÿžNoneïijŇåĹŹä¡£çŤĺéżŸèőď'åůěåĚ*fleet.util*ïijŇ éżŸèőď'äÿžNoneãĂĆ

çŤĺæşŢçd'žä¿ŃïijŽ

```
...
pred, label = model()

# 1.␣
↪å■ŢæIJžçżĎç¡ŚéŸűæőţïijŇèőąçőŮåźşæŰźèŕŕåůőåŠŇæăůæIJňæĂżæŢř
sqrerr = paddle.static.create_global_var(name="sqrerr",␣
↪persistable=True, dtype='float32', shape=[1], value=0)
total_cnt = paddle.static.create_global_var(name="total_cnt
↪", persistable=True, dtype='float32', shape=[1], value=0)

batch_cnt = paddle.sum(
    paddle.full(shape=[paddle.shape(label)[0], 1], fill_
↪value=1.0))
batch_sqrerr = paddle.nn.functional.mse_loss(pred, label,␣
↪reduction='sum')

paddle.assign(sqrerr + batch_sqrerr, sqrerr)
paddle.assign(total_cnt + batch_cnt, total_cnt)
mse =  sqrerr / total_cnt

# 2. åĹĘåÿČåijŕèő■çżČéŸűæőţïijŇèőąçőŮåĚĺåśĂåĞEçąőçŐĞãĂĆ
global_mse = fleet.metrics.mse(sqrerr, total_cnt)
```

åĹĘåÿČåijŕRMSE

paddle.distributed.fleet.metrics.**rmse**(*sqrerr*, *total_ins_num*, *scope=None*, *util=None*)

åĹĘåÿČåijŔåİĞæŰźæăźèŕŕåůőïijĹRoot Mean Squared Error)ãĂĆåİĞæŰźæăźèŕŕåůőæŸŕåİĞæŰźèŕŕåůőçŽĎçőŮæIJŕåźşæŰźæăźïijŇäžęçğřæăĞåĞĘèŕŕåůőïijŇäÿĂèĹňç*loss* æ■§åď'śåĂijãĂĆ

$$sqrerr = \sum (input - label)^2$$

$$rmse = \sqrt{\frac{sqrerr}{total\_ins\_num}}$$

åĚűäÿ■ïijŇ*input* æŸŕæăůæIJňéćĎæţŃçżŞæđIJïijŇ *label* æŸŕæăůæIJňçIJ§åődæăĞç■¿ïijŇ*sqrerr* äÿžåźşæŰźèŕŕåůőåŠŇïijŇ*total_ins_num* æŸŕæăůæIJňæĂżæŢřãĂĆ

åŔĆæŢřïijŽ

- sqrerr, (numpy.array|Tensor|string, required): å■ŢæIJžåźşæŰźèŕŕåůőåŠŇçż§èőąåĂijãĂĆ
- total_ins_num, (numpy.array|Tensor|string, required): å■ŢæIJžæăůæIJňæĂżæŢřãĂĆ
- scope, (Scope, optional)ïijŇä¡IJçŤĺå§§ïijŇèŃěäÿžNoneïijŇåĹŹä¡£çŤĺåĚĺåśĂ/éżŸèőď'ä¡IJçŤĺå§§ïijŇéż
- util, (UtilBase, optinal)ïijŇåĹĘåÿČåijŔèő■çżČåůěåĚůçśżïijŇèŃěäÿžNoneïijŇåĹŹä¡£çŤĺéŸèőď'åůěåĚ *fleet.util*ïijŇ éżŸèőď'äÿžNoneãĂĆ

çŤĺæşŢçď'žä¿ŃïjŽ

```
...
pred, label = model()

# 1.␣
↪å■ŢæIJžçżĎç¡ŚéŸűæőţïijŇèőąçőŮåźşæŰźèŕŕåůőåŠŇæăůæIJňæĂżæŢř
sqrerr = paddle.static.create_global_var(name="sqrerr",␣
↪persistable=True, dtype='float32', shape=[1], value=0)
total_cnt = paddle.static.create_global_var(name="total_cnt
↪", persistable=True, dtype='float32', shape=[1], value=0)

batch_cnt = paddle.sum(
    paddle.full(shape=[paddle.shape(label)[0], 1], fill_
↪value=1.0))
batch_sqrerr = paddle.nn.functional.mse_loss(pred, label,␣
↪reduction='sum')

paddle.assign(sqrerr + batch_sqrerr, sqrerr)
paddle.assign(total_cnt + batch_cnt, total_cnt)
mse =  sqrerr / total_cnt
rmse = paddle.sqrt(mse)

# 2. åĹĘåÿČåijŔèő■çżČéŸűæőţïijŇèőąçőŮåĚĺåśĂåĞĘçąőçŐĞãĂĆ
global_rmse = fleet.metrics.rmse(sqrerr, total_cnt)
```

åĹĘåÿČåijŔSum

paddle.distributed.fleet.metrics.**sum**(*input*, *scope=None*, *util=None*)

åĹĘåÿČåijŔæśĆåŠŇãĂĆäÿĂèĹňçŤĺäžŐèĞłåőŽäźL'æŇĞæăĞèőąçőŮãĂĆ

åŔĆæŢřïijŽ

- input, (numpy.array|Tensor|string, required)ïijŇéIJĂèęĄåĹĘåÿČåijŔæśĆåŠŇçŽĎè¿ŞåĚěåŔĆæŢřãĂĆ
- scope, (Scope, optional)ïijŇä¡IJçŤĺå§§ïijŇèŃěäÿžNoneïijŇåĹŹä¡£çŤĺåĚĺåśĂ/éżŸèőď'ä¡IJçŤĺå§§ïijŇéż
- util, (UtilBase, optinal)ïijŇåĹĘåÿČåijŔèő■çżČåůěåĚůçśżïijŇèŃěäÿžNoneïijŇåĹŹä¡£çŤĺéŸèőď'åůěåĚ *fleet.util*ïijŇ éżŸèőď'äÿžNoneãĂĆ

çŤĺæşŢçď'žä¿ŃïjŽ

```
...
# 1. å■ŢæIJžçżĎç¡ŚéŸűæőţïijŇèőąçőŮLoss
loss = model()

# 2. åĹĘåÿČåijŕèő■çżČéŸűæőţïijŇèőąçőŮåĚĺåśĂLossåŠŇ
loss_val, = exe.run(paddle.static.default_main_program(),
                    fetch_list=[loss.name])
loss_sum = fleet.metrics.sum(loss_val)
```

åĹĘåÿČåijŔMax

paddle.distributed.fleet.metrics.**max**(*input*, *scope=None*, *util=None*)

åĹĘåÿČåijŔæśĆæIJĂåď ğåĂijãĂĆäÿĂèĹňçŤĺäžŐèĞłåőŽäźĽ’æŇĞæăĞèőąçőŮãĂĆ

åŔĆæŢřïijŽ

- input, (numpy.array|Tensor|string, required)ïijŇéIJĂèęĄåĹĘåÿČåijŔæśĆæIJĂåď ğåĂijçŽĎè¿ŞåĚěåŔĆa
- scope, (Scope, optional)ïijŇä¡IJçŤĺå§§ïijŇèŃěäÿžNoneïijŇåĹŹä¡£çŤĺåĚĺåśĂ/éżŸèőď ä¡IJçŤĺå§§ïijŇéż
- util, (UtilBase, optinal)ïijŇåĹĘåÿČåijŔèő■çżČåůěåĚůçśżïijŇèŃěäÿžNoneïijŇåĹŹä¡£çŤĺéżŸèőď åůěåĚu
  *fleet.util*ïijŇ éżŸèőď äÿžNoneãĂĆ

çŤĺæşŢçď’žä¿ŃïijŽ

```
...
# 1. å■ŢæIJžçżĎç¡ŚéŸűæőţïijŇèőąçőŮLoss
loss = model()

# 2. åĹĘåÿČåijŕèő■çżČéŸűæőţïijŇèőąçőŮåĚĺåśĂæIJĂåď ğLoss
loss_val, = exe.run(paddle.static.default_main_program(),
                    fetch_list=[loss.name])
max_loss = paddle.metrics.max(loss_val)
```

åĹĘåÿČåijŔMin

paddle.distributed.fleet.metrics.**min**(*input*, *scope=None*, *util=None*)

åĹĘåÿČåijŔæśĆæIJĂåřŔåĂijãĂĆäÿĂèĹňçŤĺäžŐèĞłåőŽäźĽ’æŇĞæăĞèőąçőŮãĂĆ

åŔĆæŢřïijŽ

- input, (numpy.array|Tensor|string, required)ïijŇéIJĂèęĄåĹĘåÿČåijŔæśĆæIJĂåď ğåĂijçŽĎè¿ŞåĚěåŔĆa
- scope, (Scope, optional)ïijŇä¡IJçŤĺå§§ïijŇèŃěäÿžNoneïijŇåĹŹä¡£çŤĺåĚĺåśĂ/éżŸèőď ä¡IJçŤĺå§§ïijŇéż
- util, (UtilBase, optinal)ïijŇåĹĘåÿČåijŔèő■çżČåůěåĚůçśżïijŇèŃěäÿžNoneïijŇåĹŹä¡£çŤĺéżŸèőď åůěåĚu
  *fleet.util*ïijŇ éżŸèőď äÿžNoneãĂĆ

çŤĺæşŢçď’žä¿ŃïijŽ

```
...
# 1. å■ŢæIJžçżĎç¡ŚéŸűæőţ
loss = model()

# 2. åĹĘåÿČåijŕèő■çżČéŸűæőţïijŇèőąçőŮåĚĺåśĂæIJĂåřŔLoss
loss_val, = exe.run(paddle.static.default_main_program(),
                    fetch_list=[loss.name])
min_loss = fleet.metrics.min(loss_val)
```

7.7.4 ä¡£çŤĺæŰźæşŢ

åőŇæŢť'è£ŘèąŇçď'žä¿ŃèğĄ *examples/wide_and_deep*ãĂĆ

éĂŽè£ĞfleetrunæŇĞäžď'è£ŘèąŇåĹĘåÿČåijŕäżžåŁąãĂĆåŚ¡äžď'çď'žä¿ŃåęĆäÿŃïijŇåĚűäÿ■serve
worker_numåĹĘåĹńäÿžæIJ■åŁąèŁĆçĆźåŠŇèő■çżČèŁĆçĆźçŽĎæŢřéĞŕãĂĆ

```
fleetrun --server_num=2 --worker_num=2 train.py
```

7.8 åĹĘåÿČåijŕéćĎæţŃ

7.8.1 çőĂäżŃ

åĹĘåÿČåijŕéćĎæţŃäżžåŁąåřĘéćĎæţŃæŢřæ■őåİĞåŇĂåĹĘåÿČåijŕåIJĺåď'ŽåŕřæIJžåŹĺäÿŁïijŇæŕŕåŕřæ
all_reduce ç■L'éŻĘåŘĹéĂŽä£ąæŞ■ä¡IJåőŇæĹŘåŘĎèĞłéćĎæţŃçżŞæđIJçŽĎåŘŇæ■ěïijŇäżÕèĂŇèŐůåŔŰæ

äÿžäżĂäźĹèęĄåĄŽåĹĘåÿČåijŕéćĎæţŃïijŇéŹď'äžĘéĂŽè£ĞæŢřæ■őåźűèąŇçŽĎæŰźåijŕèŁĆçIJĄéćĎæ
çĺĂçŰŕåŔĆæŢřïijĹembeddingïijL'çŽĎ *feature id* åŕŕèČ¡äijŽéİđåÿÿåď'ŽïijŇå¡Ş *feature id*
è¿¿åĹřäÿĂåőŽæŢřéĞŕæŮűïijŇçĺĂçŰŕåŔĆæŢřäijŽåŔŸå¿Ůå¿Ĺåď'ğäžèèĞşäžŐå■ŢæIJžåĘĚå■ŸæŮăæşŢå■Ÿ

7.8.2 åŐ§çŘĘ

åĹĘåÿČåijŕéćĎæţŃçŽĎåŐ§çŘĘå§žæIJňåŠŇåĹĘåÿČåijŕèő■çżČäÿĂèĞť'ïijŇéČ¡åřĘèŁĆçĆźåĹĘäÿž
Worker åŠŇ **PServer** äÿď'çśžïijŇè£Źäÿď'çśžèŁĆçĆźåIJĺèő■çżČäżžåŁąåŠŇéćĎæţŃäżžåŁąäÿ■çŽĎåĹĘåůěåę

- **Worker**ïijŽåIJĺèő■çżČæŮűïijŇWorkerèť'§èť'čåőŇæĹŘèő■çżČæŢřæ■őèŕžåŔŰãĂĄäżŐPServeräÿŁæŃĹ
- **PServer**ïijŽåIJĺèő■çżČæŮűïijŇPServeråIJĺæŤűåĹřèő■çżČWorkeräijăæİěçŽĎæćŕåžęåŘŐïijŇäijŽæăźæ■

åĹĘåÿČåijŕéćĎæţŃäżžåŁąçŽĎæţĄçĺŃäÿżèęĄæIJL'äžěäÿŃäÿL'æ■ěïijŽ

1. èĞłåőŽäźL'éćĎæţŃçżĎç¡Ś
2. åĹİåğŃåŇŰåĹĘåÿČåijŕéŻĘç¿d'çŐŕåćČïijŇåŁăè¡¡æĺąåđŃåŔĆæŢřãĂĆ
3. çŤ§æĹŘåĹĘåÿČåijŕéćĎæţŃçżĎç¡ŚïijŇèĞłåőŽäźL'readerïijŇåijĂåğŃéćĎæţŃãĂĆ

åĹĘåÿČåijŔéćĎæţŃåŁ§èČ¡äÿżèęĄéĂŽè£Ğ *DistributedInfer* åůěåĚůçśżåőŇæĹŘïijŇäÿŃéİćåŕźçŻÿåĚşAPIçŽĎåŁ§èČ¡åŠŇåŔĆæŢřè£Żèą Ňäżńçż■ãĂĆ

class paddle.distributed.fleet.utils.ps_util.**DistributedInfer**(*main_program=*
startup_program

PaddlePaddleçŽĎåĹĘåÿČåijŔéćĎæţŃåůěåĚůçśżãĂĆ

åŔĆæŢřïijŽ

- main_program(paddle.static.Program, optional)ïijŇå■ŢæIJžéćĎæţŃçżĎç¡ŚïijŇèŃěäÿžNoneïi *paddle.static.default_main_program()* äÿžå■ŢæIJžéćĎæţŃçżĎç¡ŚãĂĆéżŸèőd'äÿžNoneãĂĆ
- startup_program(paddle.static.Program, optional)ïijŇå■ŢæIJžéćĎæţŃåĹİåğŃåŇŰçżĎç¡ŚïijŇ *paddle.static.default_startup_program()* äÿžå■ŢæIJžéćĎæţŃåĹİåğŃåŇŰçżĎç¡ŚãĂĆéżŸèőd'

æŰźæşŢïijŽ

init_distributed_infer_env(*exe*, *loss*, *role_maker=None*, *dirname=None*)

åĹİåğŃåŇŰåĹĘåÿČåijŔéŻEç¿d'çŐŕåćČïijŇåŁäè¡¡æĺąåđŃåŔĆæŢřãĂĆéIJĂèęĄæşĺæĎŔïijŇèřěæ

åŔĆæŢřïijŽ

- exe, (paddle.static.Executor, required)ïijŇåĹİåğŃåŇŰåĹĘåÿČåijŔéŻEç¿d'çŐŕåćČæŮűéI
- loss, (Tensor, required)ïijŇ écĎæţŃç¡ŚçżIJ *loss* åŔŸéĞŔãĂĆ
- role_maker, (RoleMakerBase, optional)ïijŇ åĹĘåÿČåijŔèő■çżČïijĹécĎæţŃïijL'äżżåŁąçŐŕåćČéĚ■ç¡őïijŇèŃěäÿžNoneïijŇåĹŹæąĘæđ
- dirname, (String, optional)ïijŇ åŔĆæŢřèůŕå¿ĎãĂĆèŃěäÿžNoneïijŇåĹŹäÿ■åŁäè¡¡åŔĆæŢ

get_dist_infer_program():

çŤ§æĹŘåĹĘåÿČåijŔéćĎæţŃçżĎç¡ŚãĂĆçŻÿè¿ČäžŐå■ŢæIJžéćĎæţŃçżĎç¡ŚïijŇäÿd'èĂĚåŇžåĹńä *lookup_table* çőŮå■ŘæŻ£æ■ćäÿž *distributed_lookup_table* ãĂĆ

è£ŤåŻđïijŽ ProgramïijŇåĹĘåÿČåijŔéćĎæţŃçżĎç¡ŚãĂĆ

7.8.3 ä¡£çŤĺæŰźæşŢ

åĹĘåÿČåijŔéćĎæţŃåÿÿèğĄçŽĎåžŤçŤĺåIJžæŹŕæIJL'äżěäÿŃäÿd'çğ■ïijŇåĹĘåÿČåijŔèő■çżČ+écĎæţŃäÿ

- **åĹĘåÿČåijŔèő■çżČ + écĎæţŃäÿĂä¡ŞäżżåŁą**ïijŽæŇĞåĹĘåÿČåijŔèő■çżČçżŞæİ§åŘŐïijŇWorkerèŁĆ
- **åĹĘåÿČåijŔécĎæţŃäżżåŁą**ïijŽæŇĞçžŕécĎæţŃçŽĎåĹĘåÿČåijŔäżżåŁąãĂĆè£ŹçśżäżżåŁąåIJĺè£ŻèąŇ

äÿŃéİćåĹĘåĹńäżŃçż■åŕźè£Źäÿdçğ■åĹĘåÿČåijŔécĎæţŃäżżåŁąçŽĎä¡£çŤĺæŰźæşŢïijŽ

åĹĘåÿČåijŔèő■çżČ + écĎæţŃäÿĂä¡ŞäżżåŁą

```
...
model = WideDeepModel()
model.net(is_train=True)

if fleet.is_server():
    fleet.init_server()
    fleet.run_server()
else:
    exe.run(paddle.default_startup_program())
    fleet.init_worker()

    # åĹĘåÿČåijŕèő■çżČ
    distributed_training(exe, model)

    # 1. çŤ$æĹŘå■ŢæIJžéćĎætŃçżĎç¡Ś
    test_main_program = paddle.static.Program()
    test_startup_program = paddle.static.Program()
    with paddle.static.program_guard(main_program=test_main_program,
↪ startup_program=test_startup_program):
        with paddle.utils.unique_name.guard():
            model.net(is_train=False)

    # 2.␣
↪çŤ$æĹŘåĹĘåÿČåijŕéćĎætŃçżĎç¡ŚïijŇåőŽäźL'readerïijŇè£ŻèąŇéćĎætŃ
    dist_infer = DistributedInfer(main_program=test_main_program,␣
↪startup_program=test_startup_program)
    dist_infer_program = dist_infer.get_dist_infer_program()

    test_data = WideDeepDataset(data_path="./data")
    reader = model.loader.set_sample_generator(test_data, batch_
↪size=batch_size, drop_last=True, places=place)

    reader.start()
    batch_idx = 0
    try:
        while True:
            loss_val = exe.run(program=dist_infer_program,
                               fetch_list=[model.cost.name])
            if batch_idx % 10 == 0:
                loss_val = np.mean(loss_val)
                message = "TEST ---> batch_idx: {} loss: {}\n".
↪format(batch_idx, loss_val)
    except fluid.core.EOFException:
        reader.reset()

    fleet.stop_worker()
```

åĹĘåÿČåijŕéćĎæţŃäżżåŁą

```
...

# 1. åőŽäźL'å■ŢæIJžéćĎæţŃçżĎç¡Ś
model = WideDeepModel()
model.net(is_train=False)

# 2. åĹİåğŃåŇŰåĹĘåÿČåijŕéćĎæţŃçŐŕåćČïijŇåŁäè¡¡æĺąåđŃåŔĆæŢř
dist_infer = DistributedInfer(main_program=test_main_program,␣
↪startup_program=test_startup_program)
exe = paddle.static.Executor()
dirname = "./init_params/"
dist_infer.init_distributed_infer_env(exe, model.cost,␣
↪dirname=dirname)

# 3.çŤŞæĹŘåĹĘåÿČåijŕéćĎæţŃçżĎç¡ŚïijŇåőŽäźL'readerïijŇè£Żèą
if fleet.is_worker():
    dist_infer_program = dist_infer.get_dist_infer_program()

    test_data = WideDeepDataset(data_path="./data")
    reader = model.loader.set_sample_generator(test_data, batch_
↪size=batch_size, drop_last=True, places=place)

    reader.start()
    batch_idx = 0
    try:
        while True:
            loss_val = exe.run(program=dist_infer_program,
                               fetch_list=[model.cost.name])
            if batch_idx % 10 == 0:
                loss_val = np.mean(loss_val)
                message = "TEST ---> batch_idx: {} loss: {}\n".
↪format(batch_idx, loss_val)
                print(message)
    except fluid.core.EOFException:
        reader.reset()

    fleet.stop_worker()
```

è£ŘèąŇæŰźæşŢ

åőŇæŢť'è£ŘèąŇçd'žä¿ŃèğĄ *examples/wide_and_deep*ãĂĆèŕěçd'žä¿ŃäÿžåĹĘåÿČåijŕèő■çżČ + écĎæţŃäÿĂä¡Şäżżå Łąã ĂĆ

éĚ■ç¡őåőŇæĹŘåŘŐïijŃéĂŽè£ĞfleetrunæŇĞäżd'è£ŘèąŇåĹĘåÿČåijŕäżżåŁąãĂĆåŚ¡äżd'çd'žä¿Ńåę
worker_numåĹĘåĹńäÿžæIJ■åŁąèŁĆçĆźåŠŇèő■çżČèŁĆçĆźçŽĎæŢřéĞŕãĂĆ

```
fleetrun --server_num=2 --worker_num=2 train.py
```

7.9 äžŇæňąåijĂåŕŚ

- TBA

7.10 æĮťä¡Şçďžä¿Ń

- TBA

CHAPTER 8

åď ğèğĎæĺąèŠÿéęŔ

- TBA

CHAPTER 9

èĞłçŻŚçİčèő■çżČ

- TBA

CHAPTER 10

åijźæĂğèő■çżČ

- TBA

CHAPTER 11

FleetXæĽľåśŢåůěåĚůåŇĚ

- TBA

CHAPTER 12

çŤÍæĹůFAQ

- TBA

Fleetä¡£çŤÍApache License 2.0åijĂæžŘå■Ŕèőő

I

P