How FreeBSD Boots: a soft-core MIPS perspective

Brooks Davis
SRI International

SCALE x13
February 21, 2015

Approved for public release; distribution is unlimited. This research is sponsored by the Defense Advanced Research Projects Agency (DARPA) and the Air Force Research Laboratory (AFRL), under contract FA8750-10-C-0237. The views, opinions, and/or findings contained in this article/presentation are those of the author(s)/presenter(s) and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.

UNIVERSITY OF CAMBRIDGE

Bluespec Extensible RISC Implementation

MIPS64-ISA Soft-core CPU

Multi-core & Multi-threaded

Why BERI?

spyware

spam

phishing

data

security

virus

attack

malware

virus alert!

virus detected

Banks lose over $300m

Sony hackers

80 million customer records

malware

alert!

virus detected

40 million cards compromised

Total damage to banks and retailers could exceed $18 billion.

— NYTimes


Your current options:
Redial
New Call
more

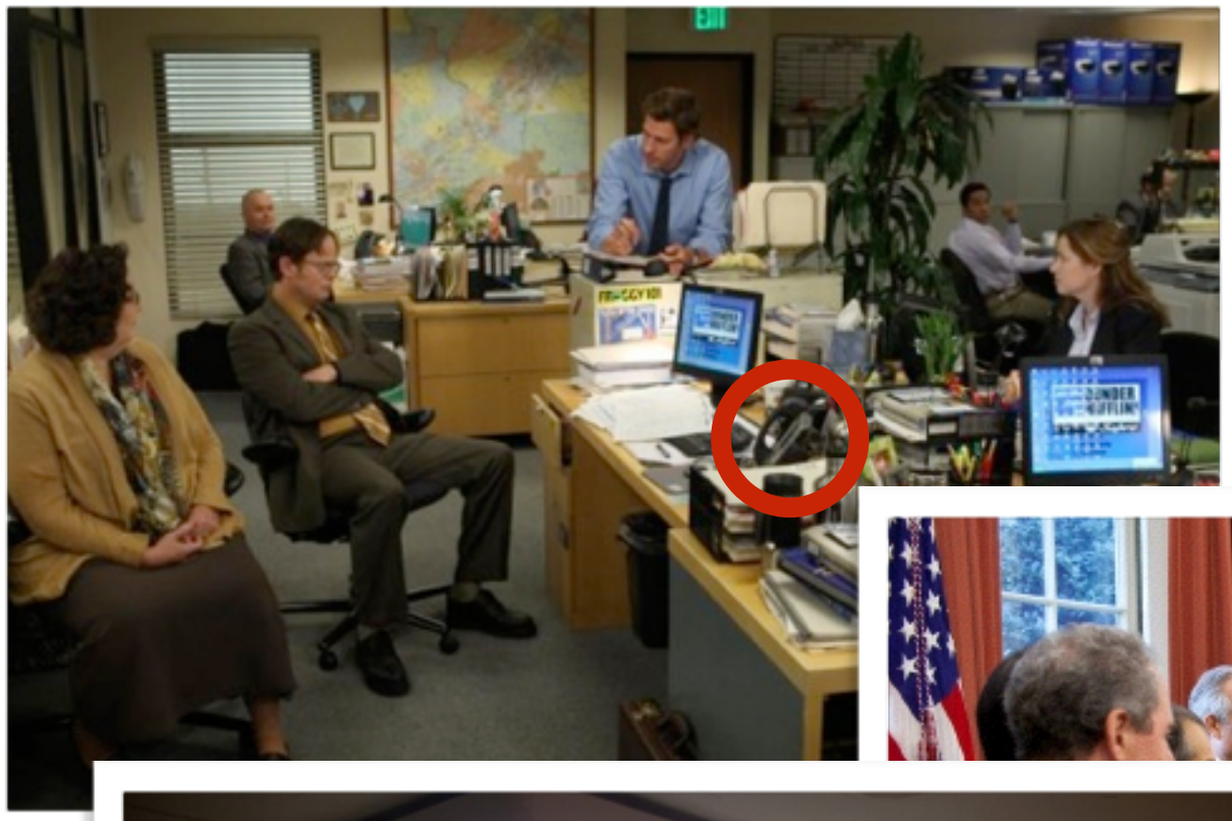


© Pete Souza | shown on planmastik.livejournal.com

CVE-2012-5445:

…allows attackers to execute arbitrary code or cause a denial of service…

(Untitled) - Wireshark

File Edit View Go Capture Analyze Statistics Help

Filter: Expression... Clear Apply

No.	Time	Source	Destination	Protocol	Info
366	11.767290	192.168.0.31	192.168.0.28	SNMP	get-response SNMPv2-SMI::enterprises.11.2.3.9.4.2.1.4.1.5.7.1
367	11.768865	192.168.0.28	192.168.0.31	SNMP	get-request SNMPv2-SMI::enterprises.11.2.3.9.4.2.1.4.1.5.8.1.
369	11.775952	192.168.0.31	192.168.0.28	SNMP	get-response SNMPv2-SMI::enterprises.11.2.3.9.4.2.1.4.1.5.8.1
381	12.286091	192.168.0.28	192.168.0.1	DNS	Standard query A www.cnn.com
384	12.311862	192.168.0.1	192.168.0.28	DNS	Standard query response A 64.236.91.21 A 64.236.91.23 A 64.23
385	12.312727	192.168.0.28	64.236.91.21	TCP	56606 > http [SYN] Seq=0 Win=8192 Len=0 MSS=1460 WS=2
386	12.361495	64.236.91.21	192.168.0.28	TCP	http > 56606 [SYN, ACK] Seq=0 Ack=1 Win=8192 Len=0 MSS=1460
387	12.361583	192.168.0.28	64.236.91.21	TCP	56606 > http [ACK] Seq=1 Ack=1 Win=17520 Len=0
388	12.361805	192.168.0.28	64.236.91.21	HTTP	GET / HTTP/1.1
389	12.413166	64.236.91.21	192.168.0.28	TCP	http > 56606 [ACK] Seq=1 Ack=845 Win=6960 Len=0
390	12.413611	64.236.91.21	192.168.0.28	TCP	[TCP segment of a reassembled PDU]
391	12.414386	64.236.91.21	192.168.0.28	TCP	[TCP segment of a reassembled PDU]

```
⊞ Frame 384 (167 bytes on wire, 167 bytes captured)
⊞ Ethernet II, Src: Sparklan_04:d0:9e (00:0e:8e:04:d0:9e), Dst: HonHaiPr_26:66:a2 (00:1c:26:26:66:a2)
⊞ Internet Protocol, Src: 192.168.0.1 (192.168.0.1), Dst: 192.168.0.28 (192.168.0.28)
⊞ User Datagram Protocol, Src Port: domain (53), Dst Port: 62872 (62872)
⊟ Domain Name System (response)
    [Request In: 381]
    [Time: 0.025771000 seconds]
    Transaction ID: 0xcf1f
  ⊞ Flags: 0x8180 (Standard query response, No error)
    Questions: 1
    Answer RRs: 6
    Authority RRs: 0
    Additional RRs: 0
  ⊟ Queries
    ⊟ www.cnn.com: type A, class IN
        Name: www.cnn.com
        Type: A (Host address)
        Class: IN (0x0001)
  ⊟ Answers
    ⊞ www.cnn.com: type A, class IN, addr 64.236.91.21
```

```
0000  00 1c 26 26 66 a2 00 0e  8e 04 d0 9e 08 00 45 00   ..&&f... ......E.
0010  00 99 00 00 40 00 40 11  b8 e6 c0 a8 00 01 c0 a8   ....@.@. ........
0020  00 1c 00 35 f5 98 00 85  98 5a cf 1f 81 80 00 01   ...5.... .Z......
0030  00 06 00 00 00 00 03 77  77 77 03 63 6e 6e 03 63   .......w ww.cnn.c
0040  6f 6d 00 00 01 00 01 c0  0c 00 01 00 01 00 00 00   om...... ........
0050  b7 00 04 40 ec 5b 15 c0  0c 00 01 00 01 00 00 00   ...@.[.. ........
0060  b7 00 04 40 ec 5b 17 c0  0c 00 01 00 01 00 00 00   ...@.[.. ........
0070  b7 00 04 40 ec 10 14 c0  0c 00 01 00 01 00 00 00   ...@.... ........
```

This is a response to the DNS query in this fr... | Packets: 1273 Displayed: 909 Marked: 0 Dropped: 0 | Profile: Default

(Untitled) - Wireshark

File Edit View Go Capture Analyze Statistics Help

Filter: Expression... Clear Apply

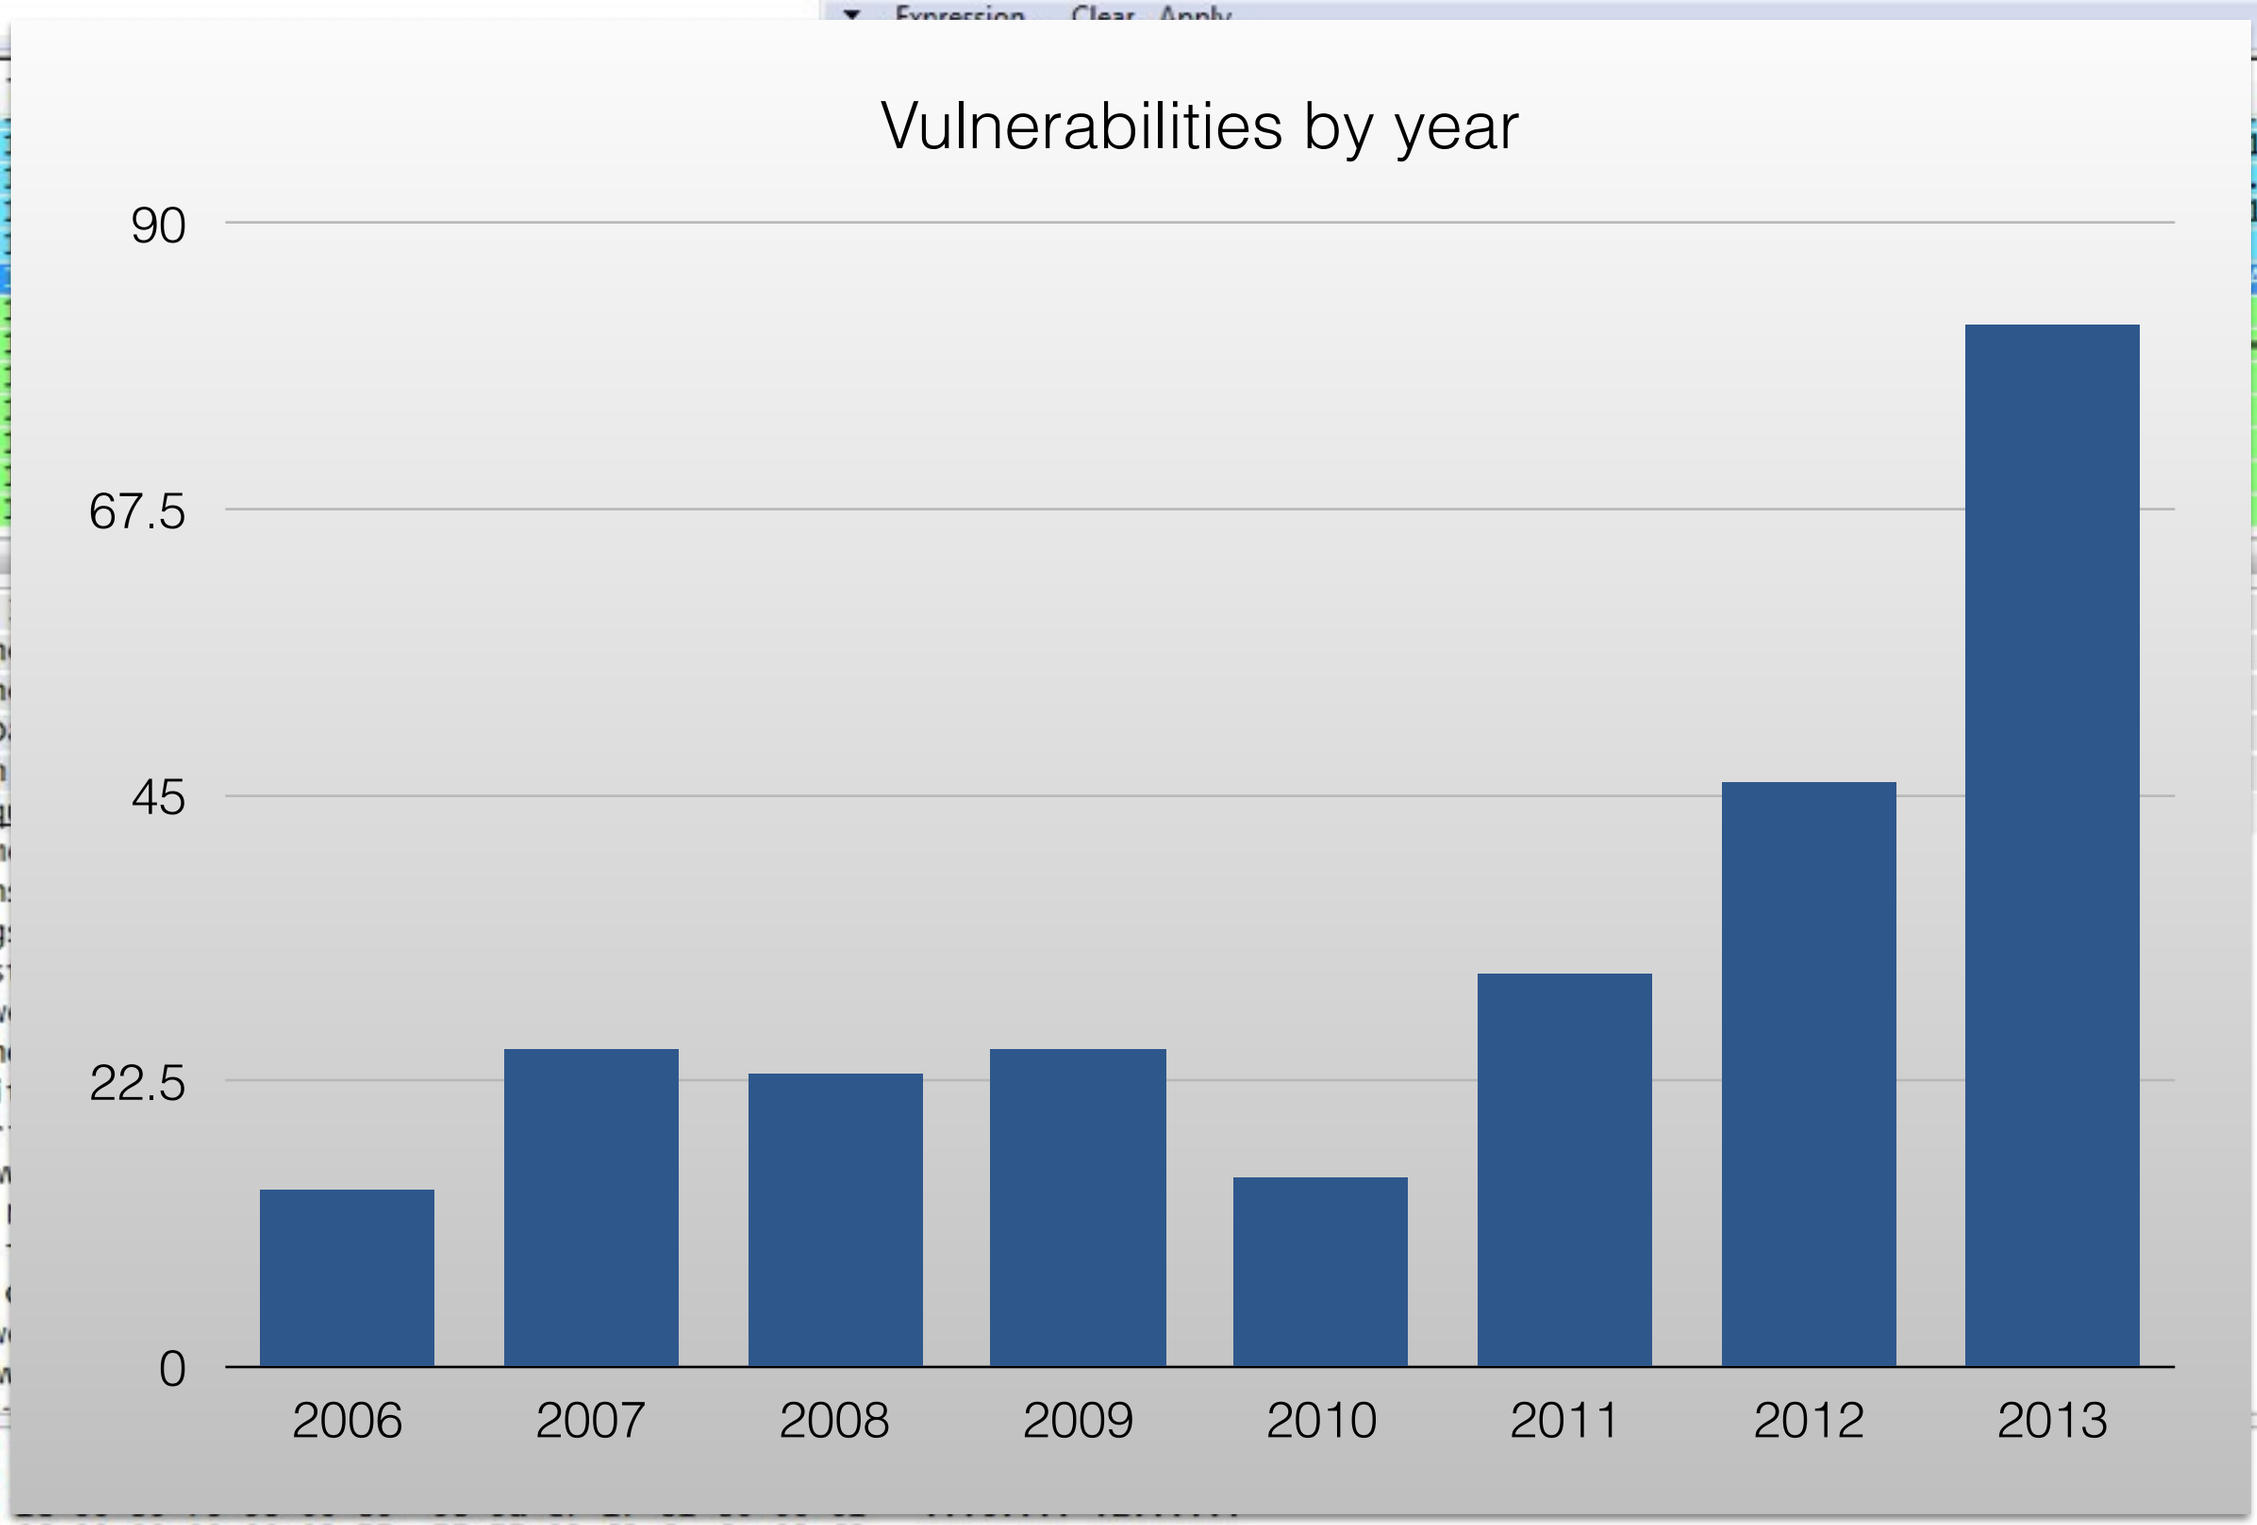

No. .

Frame
Ethern
Intern
User D
Domain
[Requ
[Time
Trans
Flags
Quest
Answe
Authe
Addit
Quer
ww
Answe
ww

0000 00
0010 00
0020 00
0030 00 06 00 00 00 00 03 77 77 77 03 63 6e 6e 03 63 .....w ww.cnn.c
0040 6f 6d 00 00 01 00 01 c0 0c 00 01 00 01 00 00 00 om..... ....
0050 b7 00 04 40 ec 5b 15 c0 0c 00 01 00 01 00 00 00 ...@.[.. ....
0060 b7 00 04 40 ec 5b 17 c0 0c 00 01 00 01 00 00 00 ...@.[.. ....
0070 b7 00 04 40 ec 10 14 c0 0c 00 01 00 01 00 00 00 ...@.... ....

This is a response to the DNS query in this fr... | Packets: 1273 Displayed: 909 Marked: 0 Dropped: 0 | Profile: Default

Clean-slate design of
Resilient,
Adaptive,
Secure
Hosts

The FreeBSD Project | Computer Laboratory: CTS | Computer Laboratory - CT | BERI Open Systems C.I.C. | YouTube

www.freebsd.org

FreeBSD® The Power To Serve

Donate to FreeBSD

Search | Search

Home | About | Get FreeBSD | Documentation | Community | Developers | Support | Foundation

de en es fr hu it ja nl ru zh_CN

IPv6 Armenia | Go

The FreeBSD Project

FreeBSD is an advanced computer operating system used to power modern servers, desktops and embedded platforms. A large community has continually developed it for more than thirty years. Its advanced networking, security and storage features have made FreeBSD the platform of choice for many of the busiest web sites and most pervasive embedded networking and storage devices.

» Learn More

Download FreeBSD

LATEST RELEASES

» Production: 10.0
» Legacy: 9.2, 8.4

SHORTCUTS

» Mailing Lists
» Report a Bug
» FAQ
» Handbook
» Ports

New to FreeBSD?

LATEST NEWS

2014-02-10
FreeBSD Journal First Edition Available

2014-01-25
October-December, 2013 Status Report

2014-01-23
New committer: Rodrigo Osorio (ports)

2014-01-21
New committer: Tycho Nightingale (src)

2014-01-21
New committer: Michael Gmelin (ports)

More News RSS

UPCOMING EVENTS

2014-03-12 - 2014-03-12
bhyvecon 2014 (Tokyo, Japan)

2014-03-13 - 2014-03-16
AsiaBSDCon 2014 (Tokyo, Japan)

2014-05-14 - 2014-05-17
BSDCan 2014 (Ottawa, Canada)

More Events

PRESS

2014-01
McKusick Denies FreeBSD Lagging on Security

2013-06
Fixing Network Attached Storage with commodity hardware and BSD

2013-02
2012 - A BSD Year in Retrospective

2013-01
A Decade of OS Access-control Extensibility

2012-11
A world without Linux: Where would Apache, Microsoft — even Apple be today?

More Media

SECURITY ADVISORIES

2014-01-14
FreeBSD-SA-14:04.bind

2014-01-14
FreeBSD-SA-14:03.openssl

2014-01-14
FreeBSD-SA-14:02.ntpd

2014-01-14
FreeBSD-SA-14:01.bsnmpd

More RSS

ERRATA NOTICES

2014-01-14
FreeBSD-EN-14:02.mmap

2014-01-14
FreeBSD-EN-14:01.random

More RSS

The FreeBSD Project | Computer Laboratory: CTS | Computer Laboratory - CT | BERI Open Systems C.I.C. | YouTube

www.freebsd.org

FreeBSD® The Power To Serve

Donate to FreeBSD

Search | Search

Home | About | Get FreeBSD | Documentation | Community | Developers | Support | Foundation

de en es fr hu it ja nl ru zh_CN

IPv6 Armenia | Go

The FreeBSD Project

FreeBSD is an advanced computer operating system used to power modern servers, desktops and embedded platforms. A large community has continually developed it for more than thirty years. Its advanced networking, security and storage features have made FreeBSD the platform of choice for many of the busiest web sites and most pervasive embedded networking and storage devices.

» Learn More

Download FreeBSD

LATEST RELEASES

» Production: 10.0
» Legacy: 9.2, 8.4

SHORTCUTS

» Mailing Lists
» Report a Bug
» FAQ
» Handbook
» Ports

New to FreeBSD?

LATEST NEWS

2014-02-10
FreeBSD Journal First Edition Available

2014-01-25
October-December, 2013 Status Report

2014-01-23
New committer: Rodrigo Osorio (ports)

2014-01-21
New committer: Tycho Nightingale (src)

2014-01-21
New committer: Michael Gmelin (ports)

More News RSS

UPCOMING EVENTS

2014-03-12 - 2014-03-12
bhyvecon 2014 (Tokyo, Japan)

2014-03-13 - 2014-03-16
AsiaBSDCon 2014 (Tokyo, Japan)

2014-05-14 - 2014-05-17
BSDCan 2014 (Ottawa, Canada)

More Events

PRESS

2014-01
McKusick Denies FreeBSD Lagging on Security

2013-06
Fixing Network Attached Storage with commodity hardware and BSD

2013-02
2012 - A BSD Year in Retrospective

2013-01
A Decade of OS Access-control Extensibility

2012-11
A world without Linux: Where would Apache, Microsoft — even Apple be today?

More Media

SECURITY ADVISORIES

2014-01-14
FreeBSD-SA-14:04.bind

2014-01-14
FreeBSD-SA-14:03.openssl

2014-01-14
FreeBSD-SA-14:02.ntpd

2014-01-14
FreeBSD-SA-14:01.bsnmpd

More RSS

ERRATA NOTICES

2014-01-14
FreeBSD-EN-14:02.mmap

2014-01-14
FreeBSD-EN-14:01.random

More RSS

Capability
Hardware
Enhanced
RISC
Instructions

BERI

+ Memory Capabilities

=

CHERI

Memory Capabilities

- Unforgeable references to memory
- Implemented as a MIPS coprocessor
- All memory access via a capability
  - Regular MIPS instructions access via a default capability

Other BERI research topics

- Novel multi-threaded CPUs
- Multi-core architectures
- Tight integration of CPUs and switches
- Validation of formal models

Where does BERI run?

Simulation

NetFPGA-10G

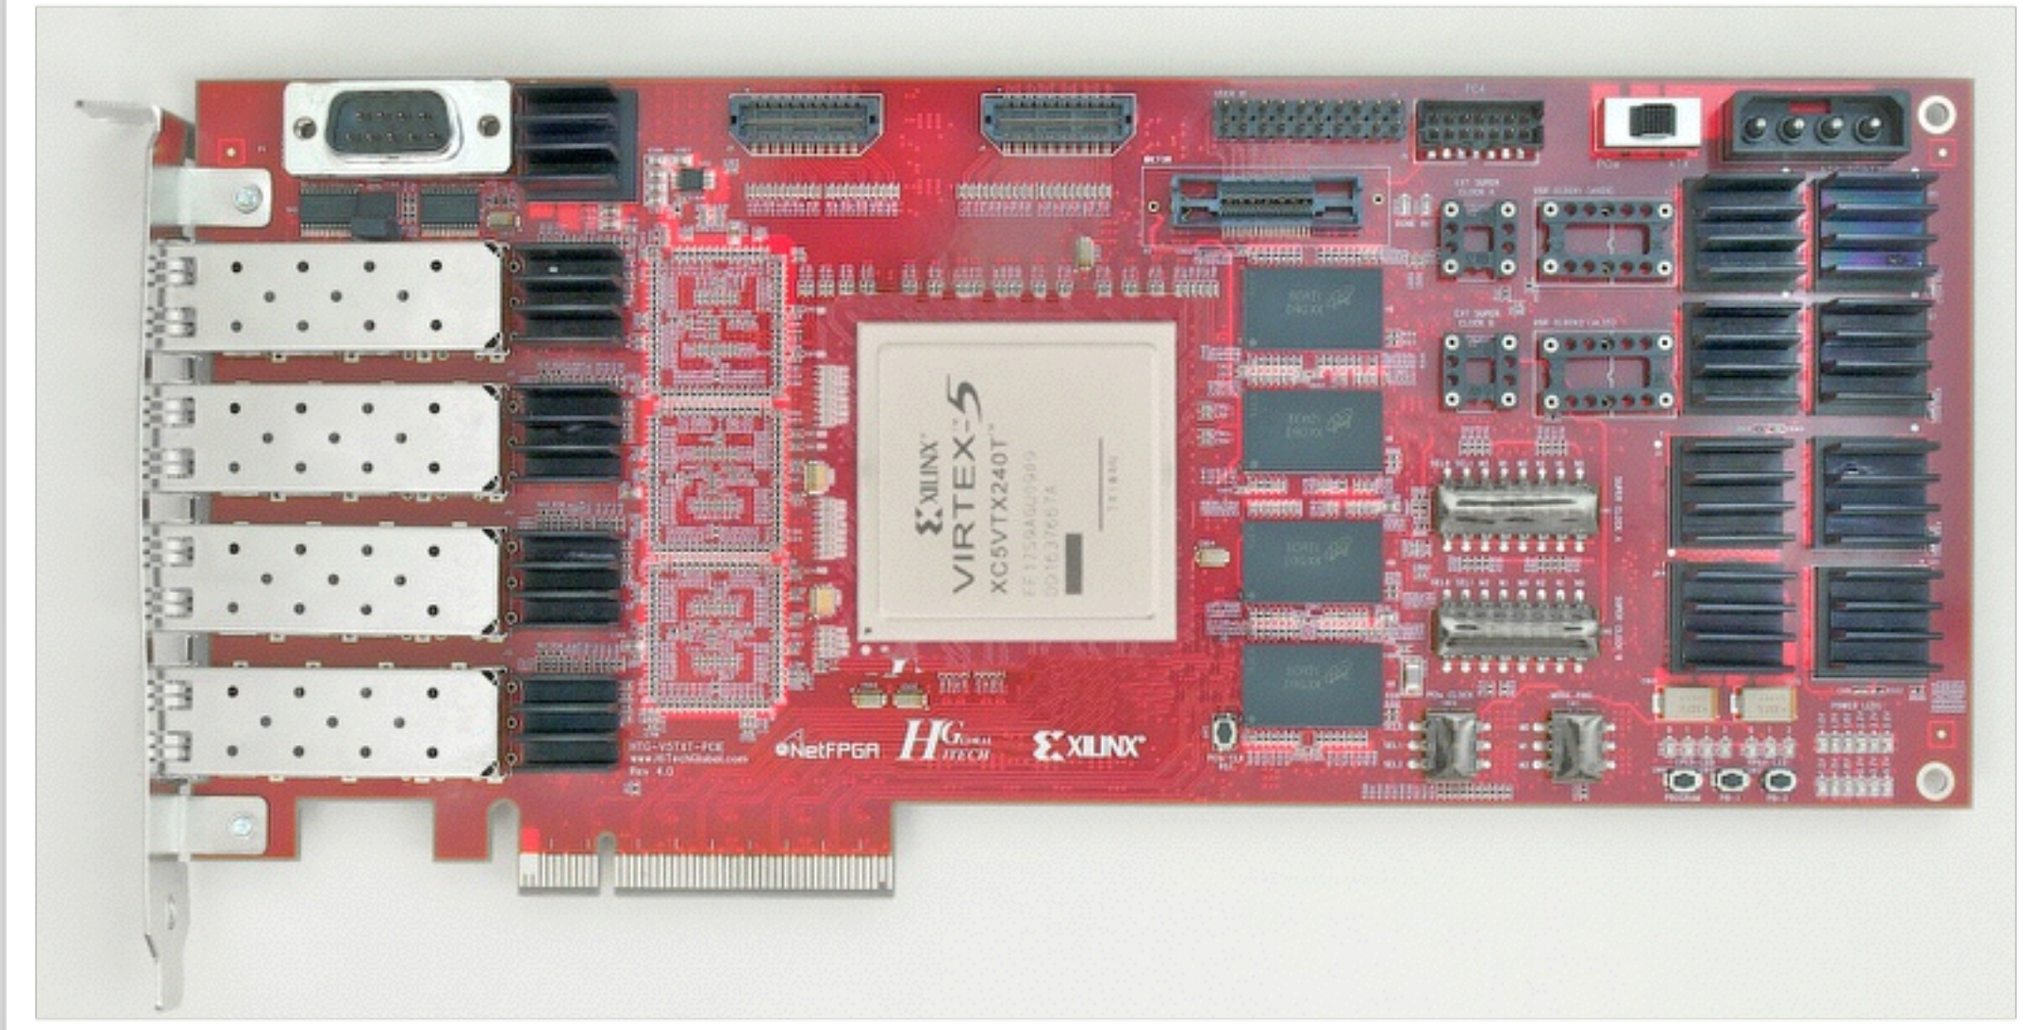

Terasic SoCKIT

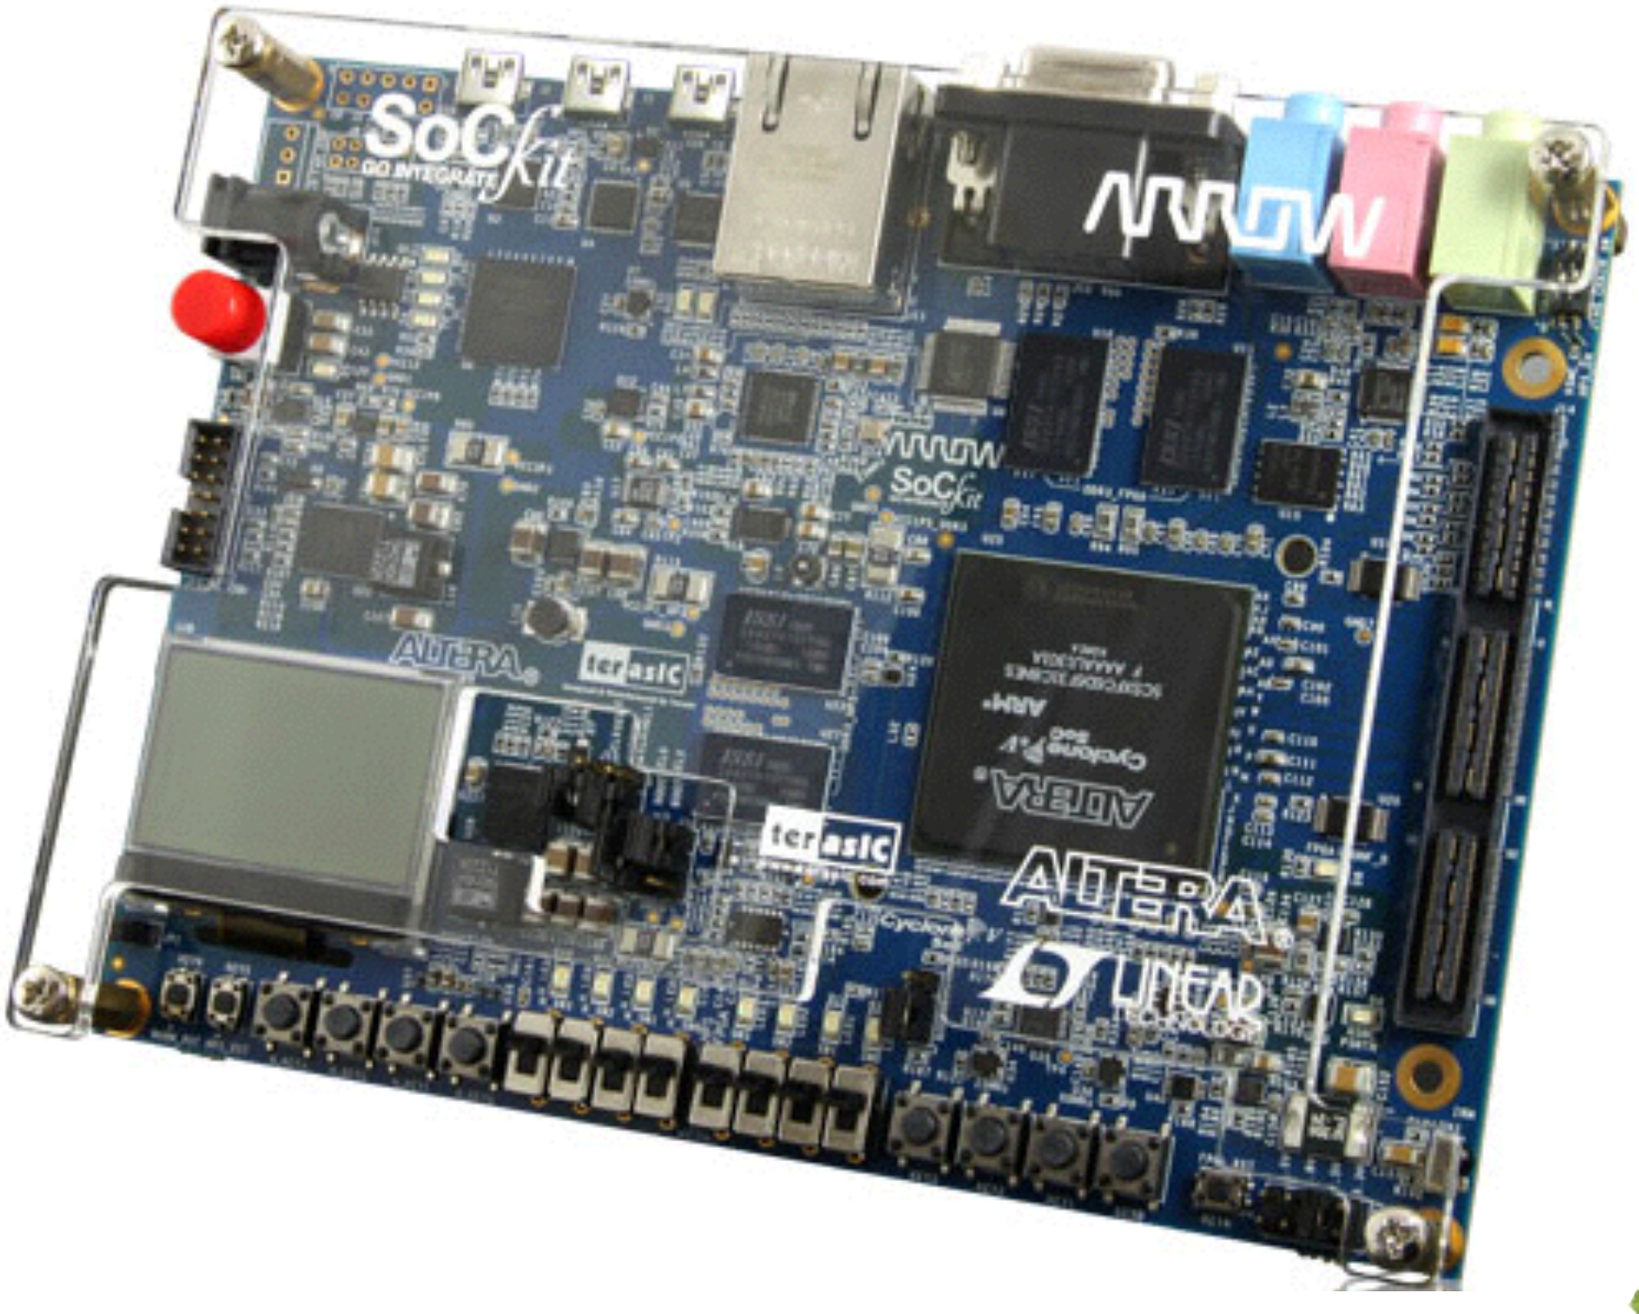

Terasic DE4

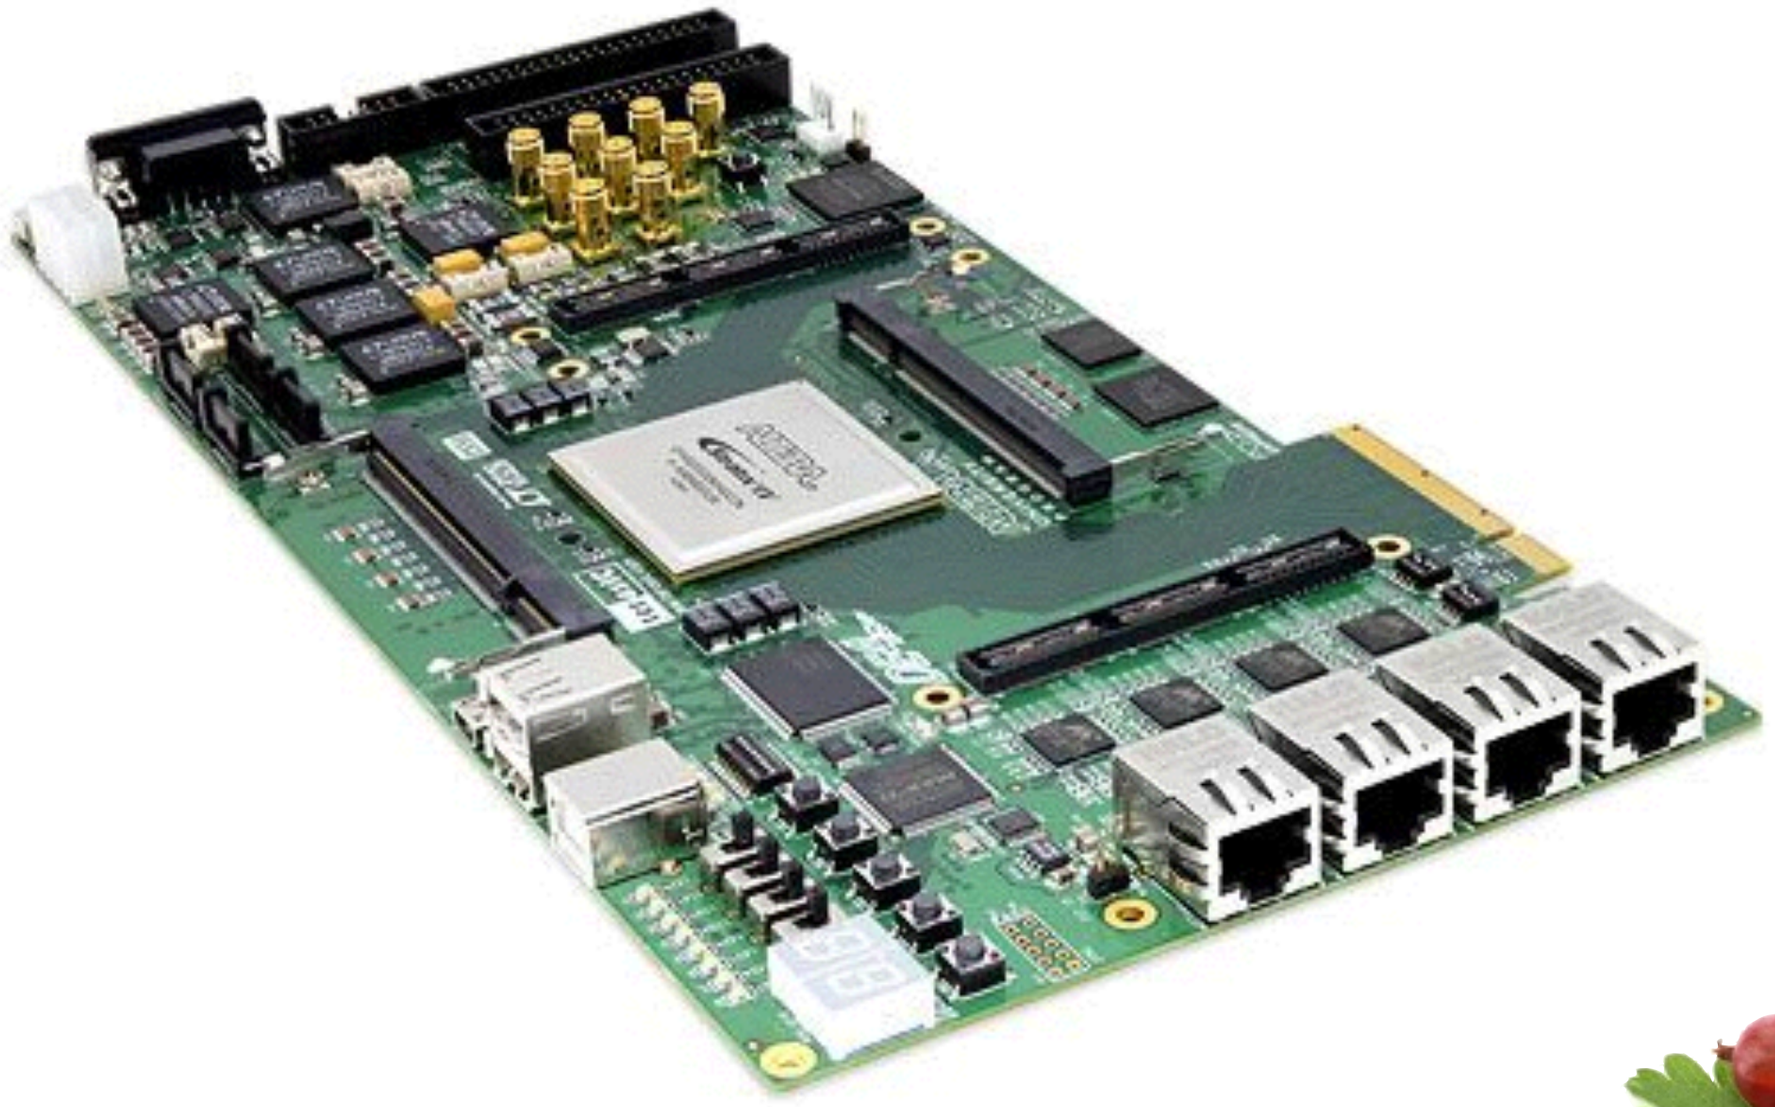


terasIC


Energizer
energi TO GO.

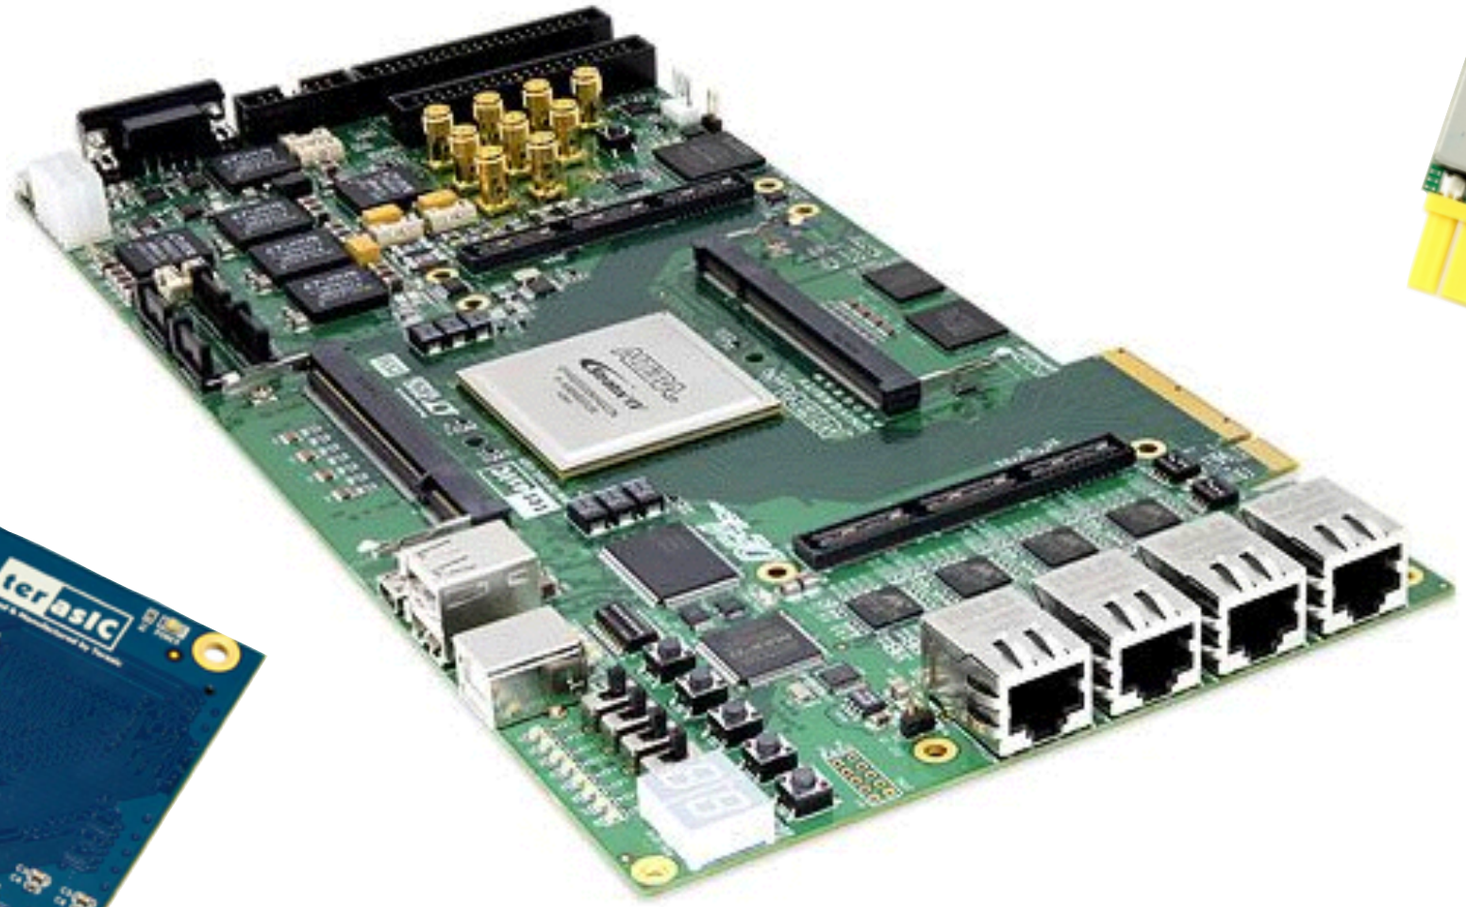


terasIC
HDMI_TX_INSMC


CTSRD
UNIVERSITY OF CAMBRIDGE
CTSRD
On-Tablet Briefing
CTSRD
Demonstration
CheriBrowser
UNIVERSITY OF CAMBRIDGE

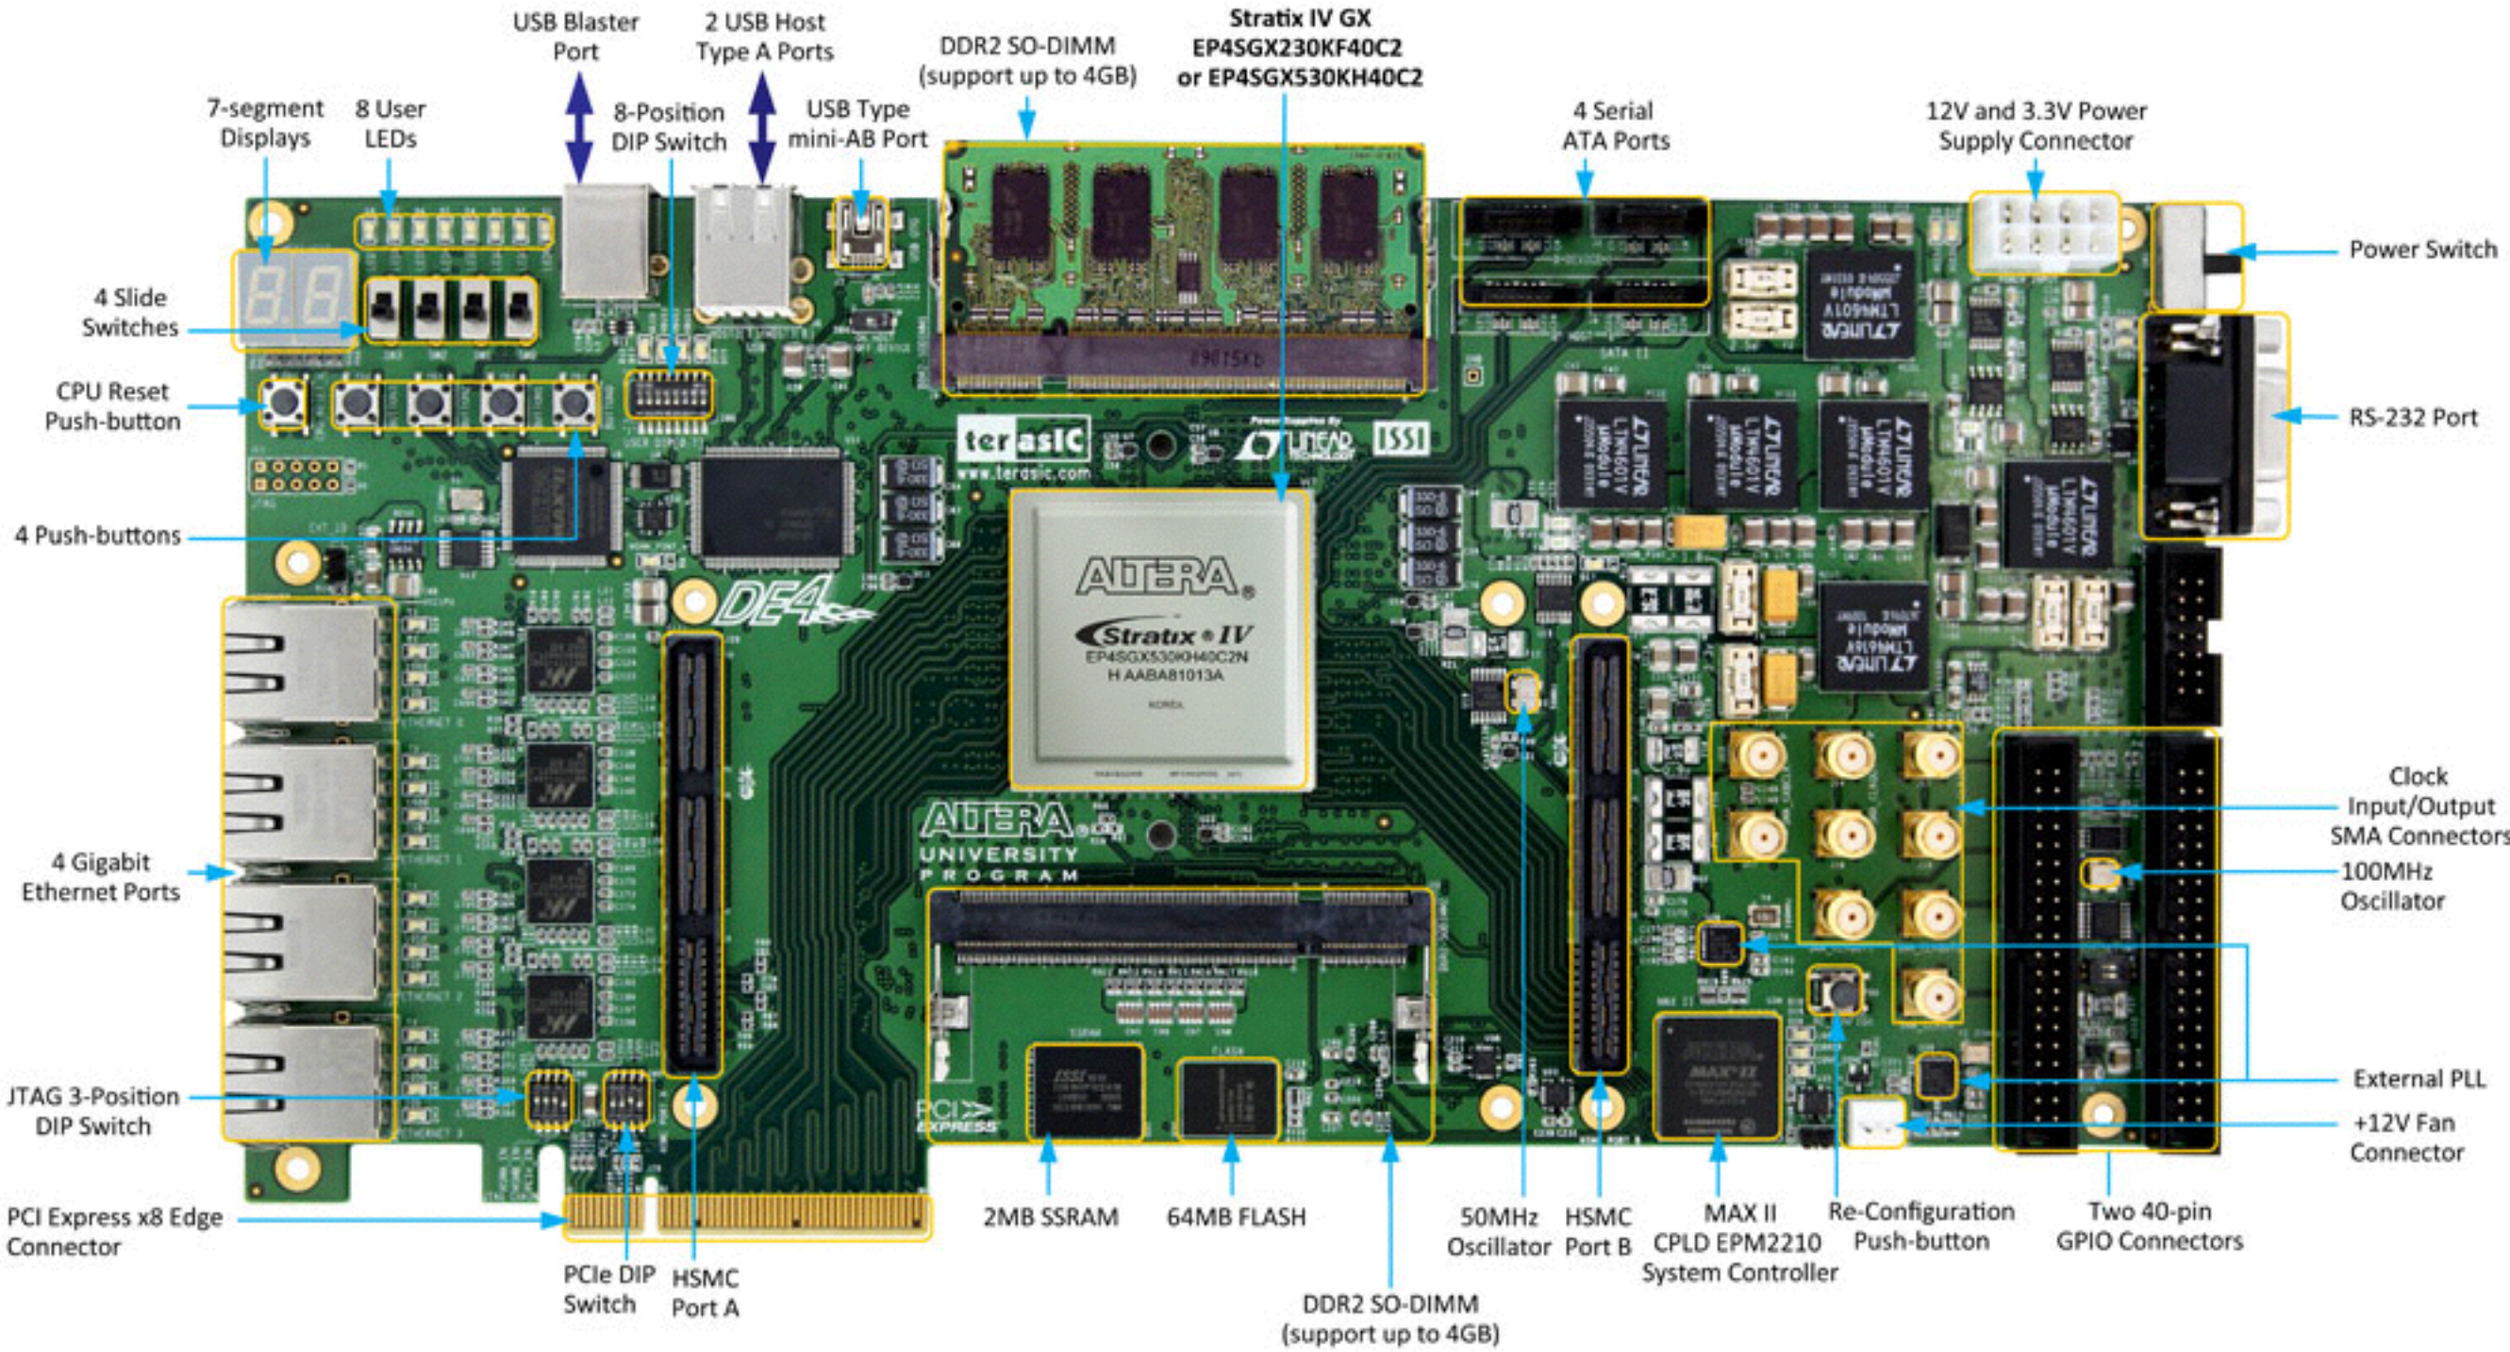
USB Blaster Port
2 USB Host Type A Ports
DDR2 SO-DIMM (support up to 4GB)
Stratix IV GX EP4SGX230KF40C2 or EP4SGX530KH40C2
7-segment Displays
8 User LEDs
8-Position DIP Switch
USB Type mini-AB Port
4 Serial ATA Ports
12V and 3.3V Power Supply Connector
Power Switch
4 Slide Switches
CPU Reset Push-button
RS-232 Port
4 Push-buttons
Clock Input/Output SMA Connectors
100MHz Oscillator
4 Gigabit Ethernet Ports
External PLL
+12V Fan Connector
JTAG 3-Position DIP Switch
PCI Express x8 Edge Connector
2MB SSRAM
64MB FLASH
50MHz Oscillator
HSMC Port B
MAX II CPLD EPM2210 System Controller
Re-Configuration Push-button
Two 40-pin GPIO Connectors
PCIe DIP Switch
HSMC Port A
DDR2 SO-DIMM (support up to 4GB)

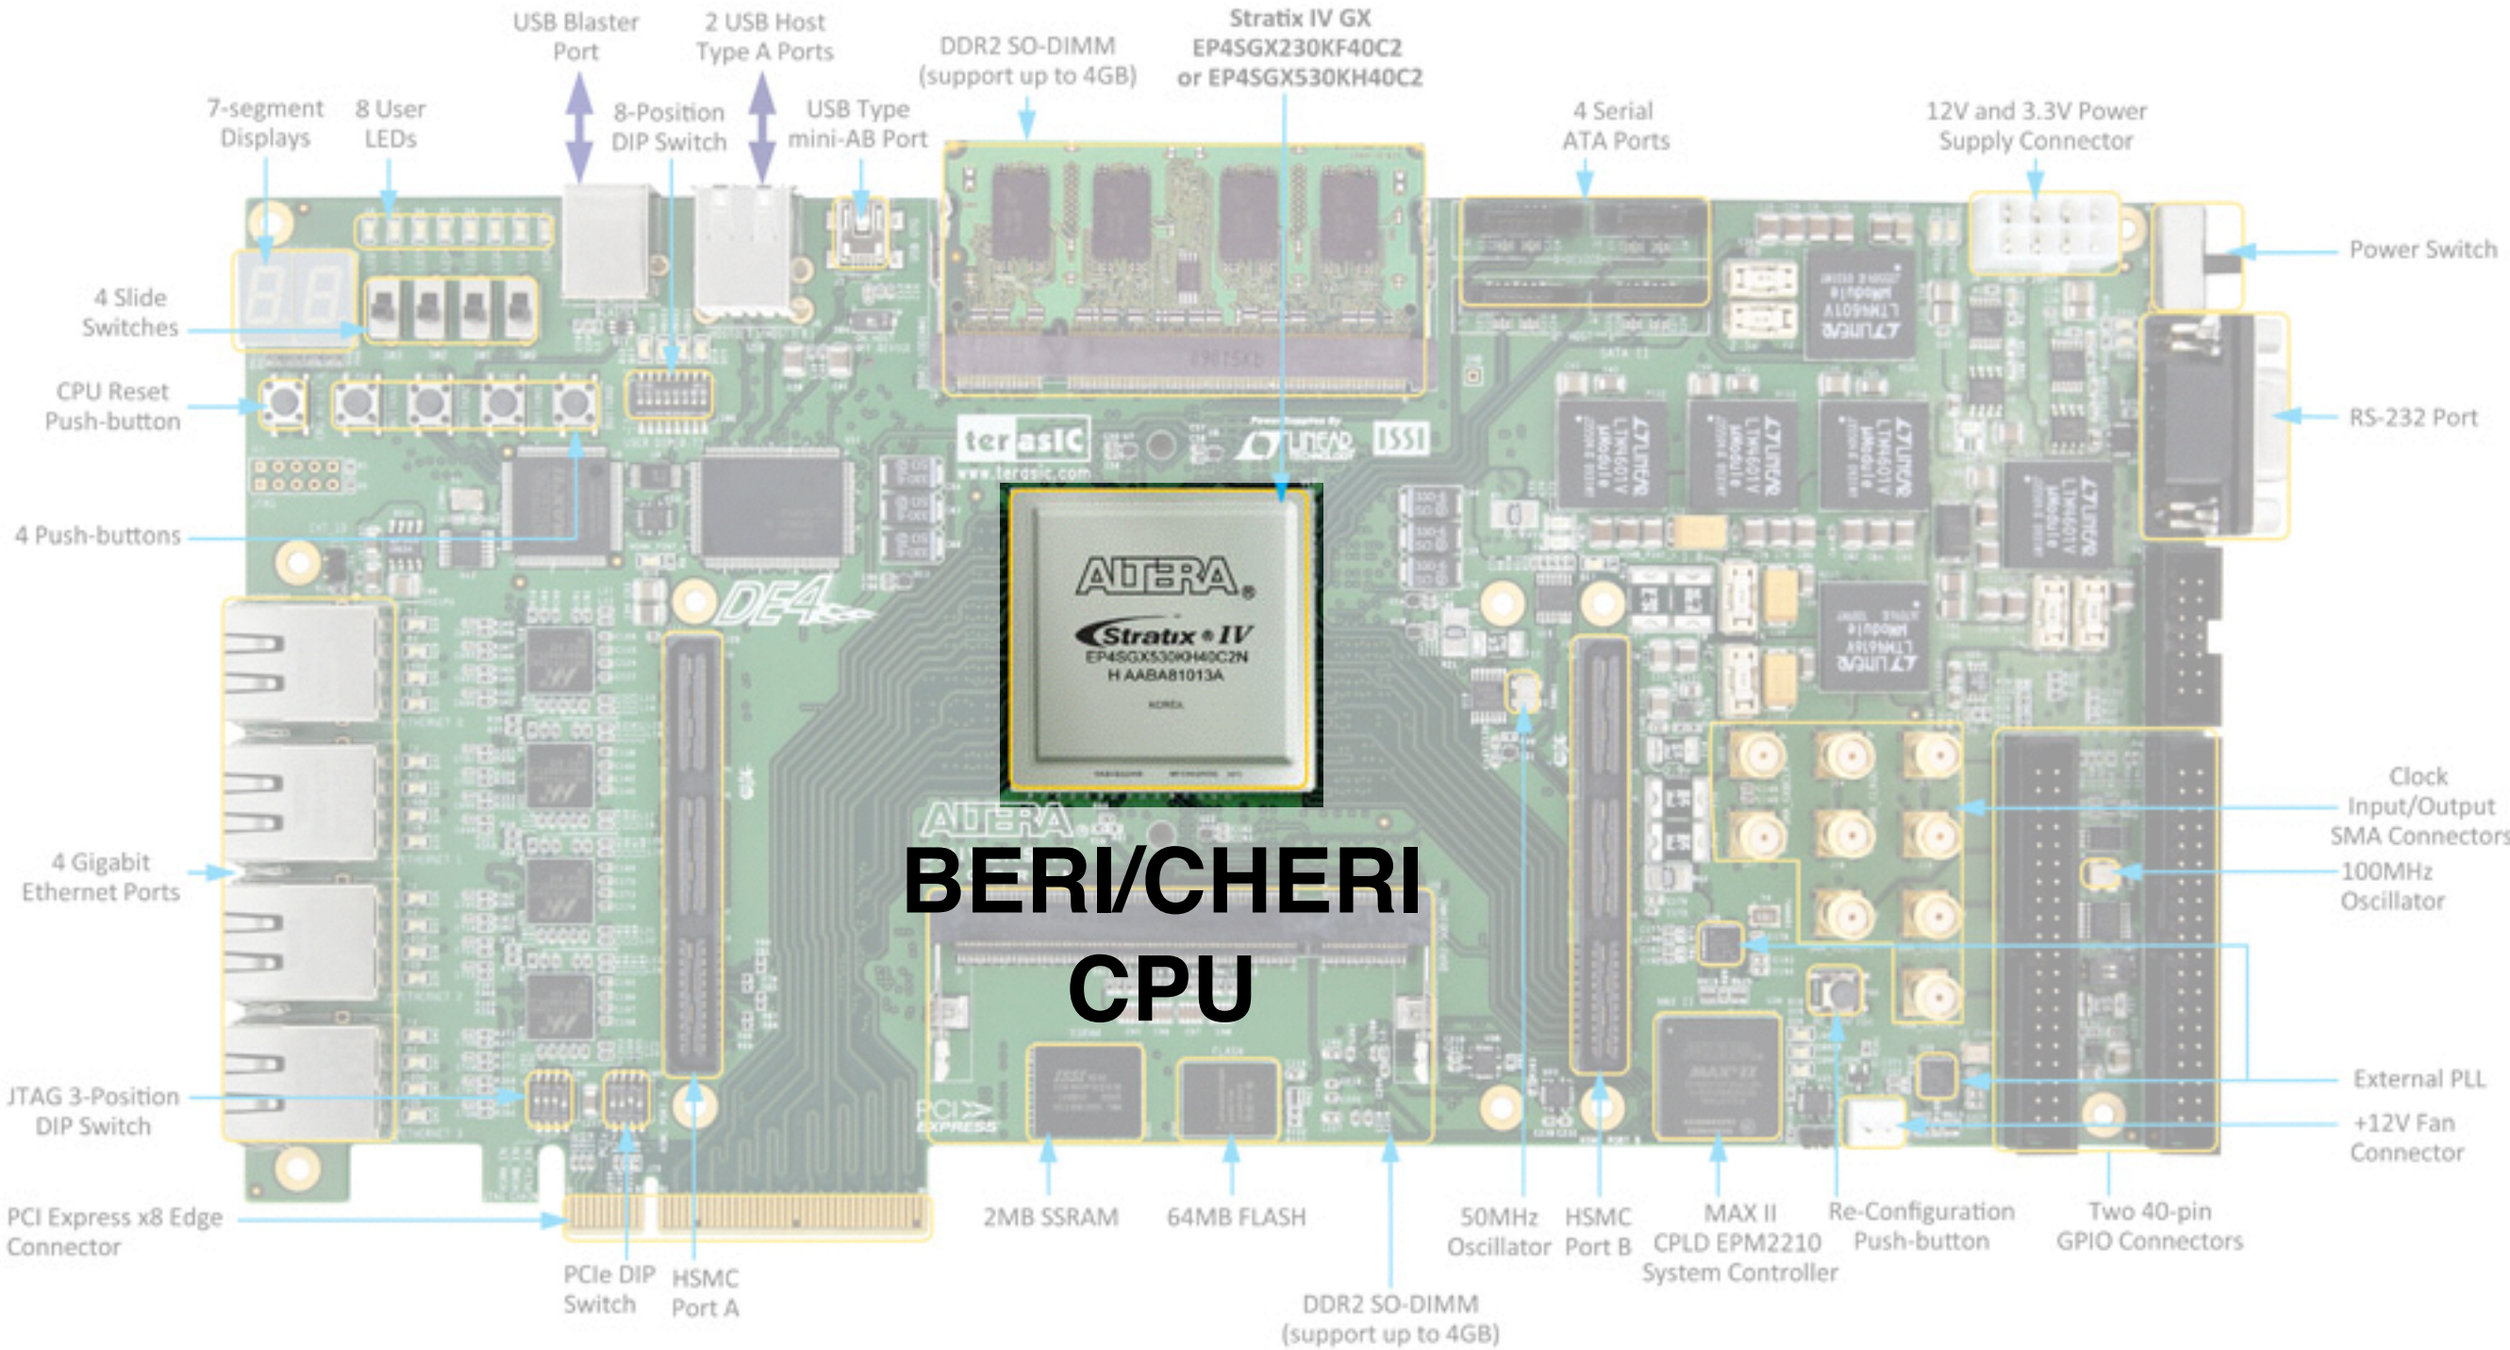
USB Blaster Port
2 USB Host Type A Ports
DDR2 SO-DIMM (support up to 4GB)
Stratix IV GX EP4SGX230KF40C2 or EP4SGX530KH40C2
7-segment Displays
8 User LEDs
8-Position DIP Switch
USB Type mini-AB Port
4 Serial ATA Ports
12V and 3.3V Power Supply Connector
Power Switch
4 Slide Switches
CPU Reset Push-button
RS-232 Port
4 Push-buttons
BERI/CHERI CPU
4 Gigabit Ethernet Ports
Clock Input/Output SMA Connectors
100MHz Oscillator
JTAG 3-Position DIP Switch
External PLL
+12V Fan Connector
PCI Express x8 Edge Connector
2MB SSRAM
64MB FLASH
50MHz Oscillator
HSMC Port B
MAX II CPLD EPM2210 System Controller
Re-Configuration Push-button
Two 40-pin GPIO Connectors
PCIe DIP Switch
HSMC Port A
DDR2 SO-DIMM (support up to 4GB)

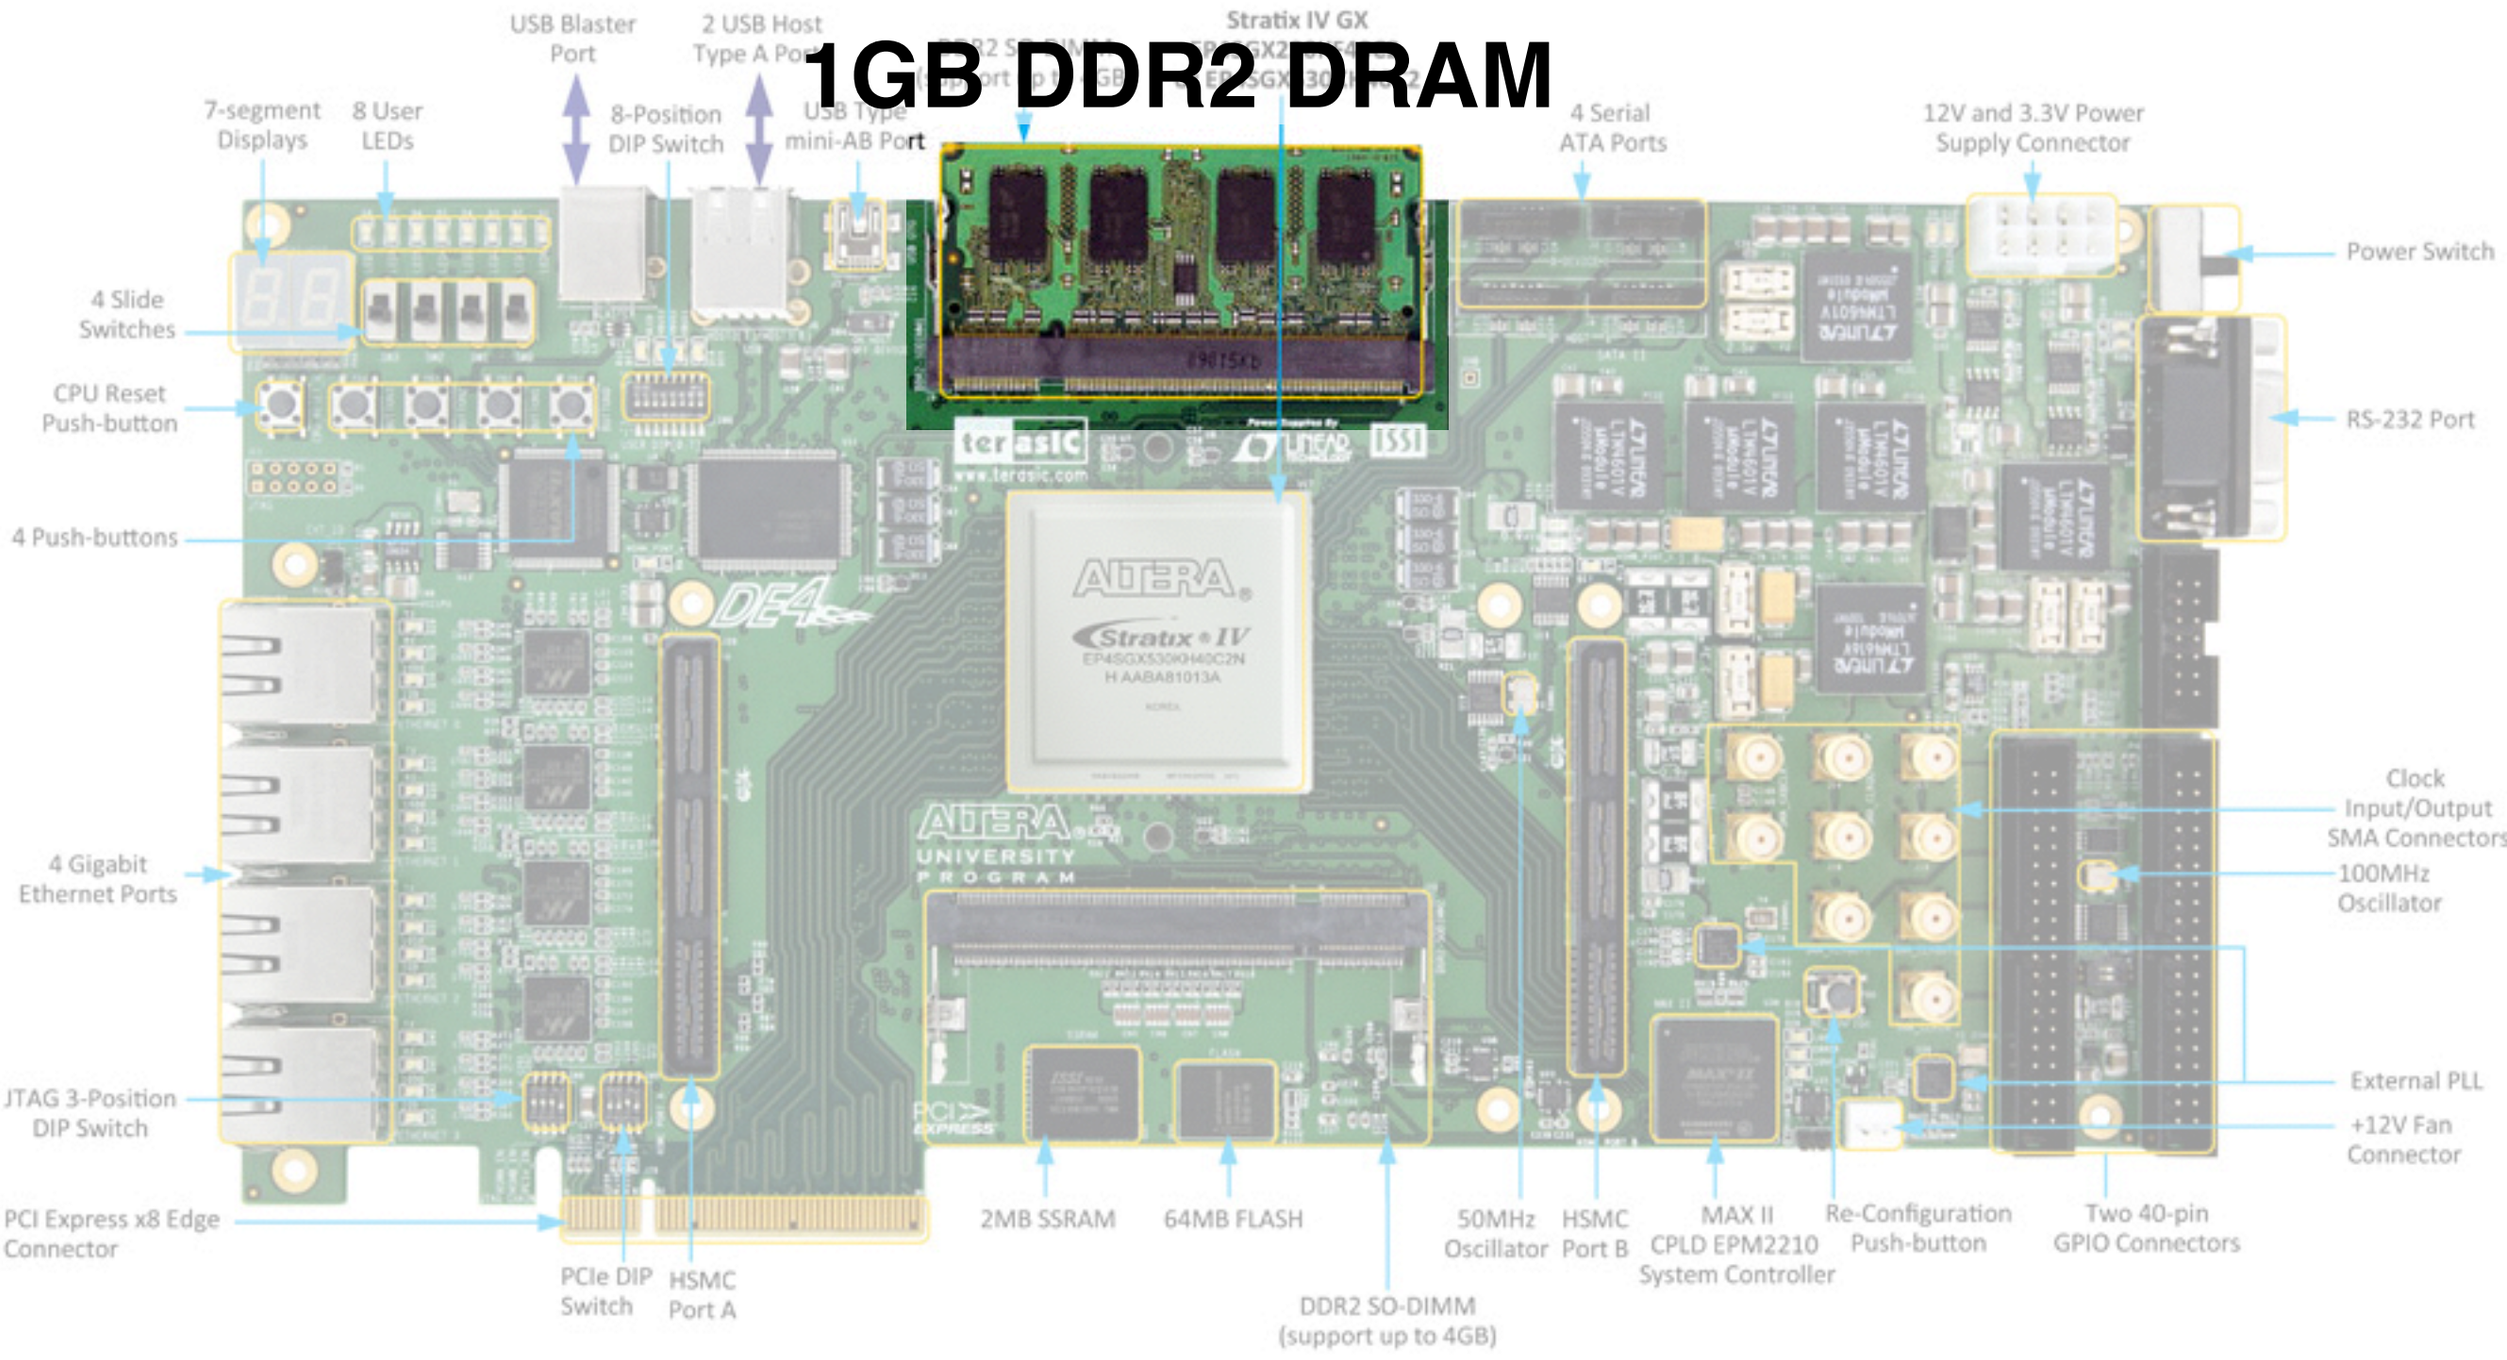
1GB DDR2 DRAM
USB Blaster Port
2 USB Host Type A Port
USB Type mini-AB Port
Stratix IV GX
7-segment Displays
8 User LEDs
8-Position DIP Switch
4 Serial ATA Ports
12V and 3.3V Power Supply Connector
Power Switch
4 Slide Switches
CPU Reset Push-button
RS-232 Port
4 Push-buttons
4 Gigabit Ethernet Ports
Clock Input/Output SMA Connectors
100MHz Oscillator
JTAG 3-Position DIP Switch
External PLL
+12V Fan Connector
PCI Express x8 Edge Connector
2MB SSRAM
64MB FLASH
50MHz Oscillator
HSMC Port B
MAX II CPLD EPM2210 System Controller
Re-Configuration Push-button
Two 40-pin GPIO Connectors
PCIe DIP Switch
HSMC Port A
DDR2 SO-DIMM (support up to 4GB)

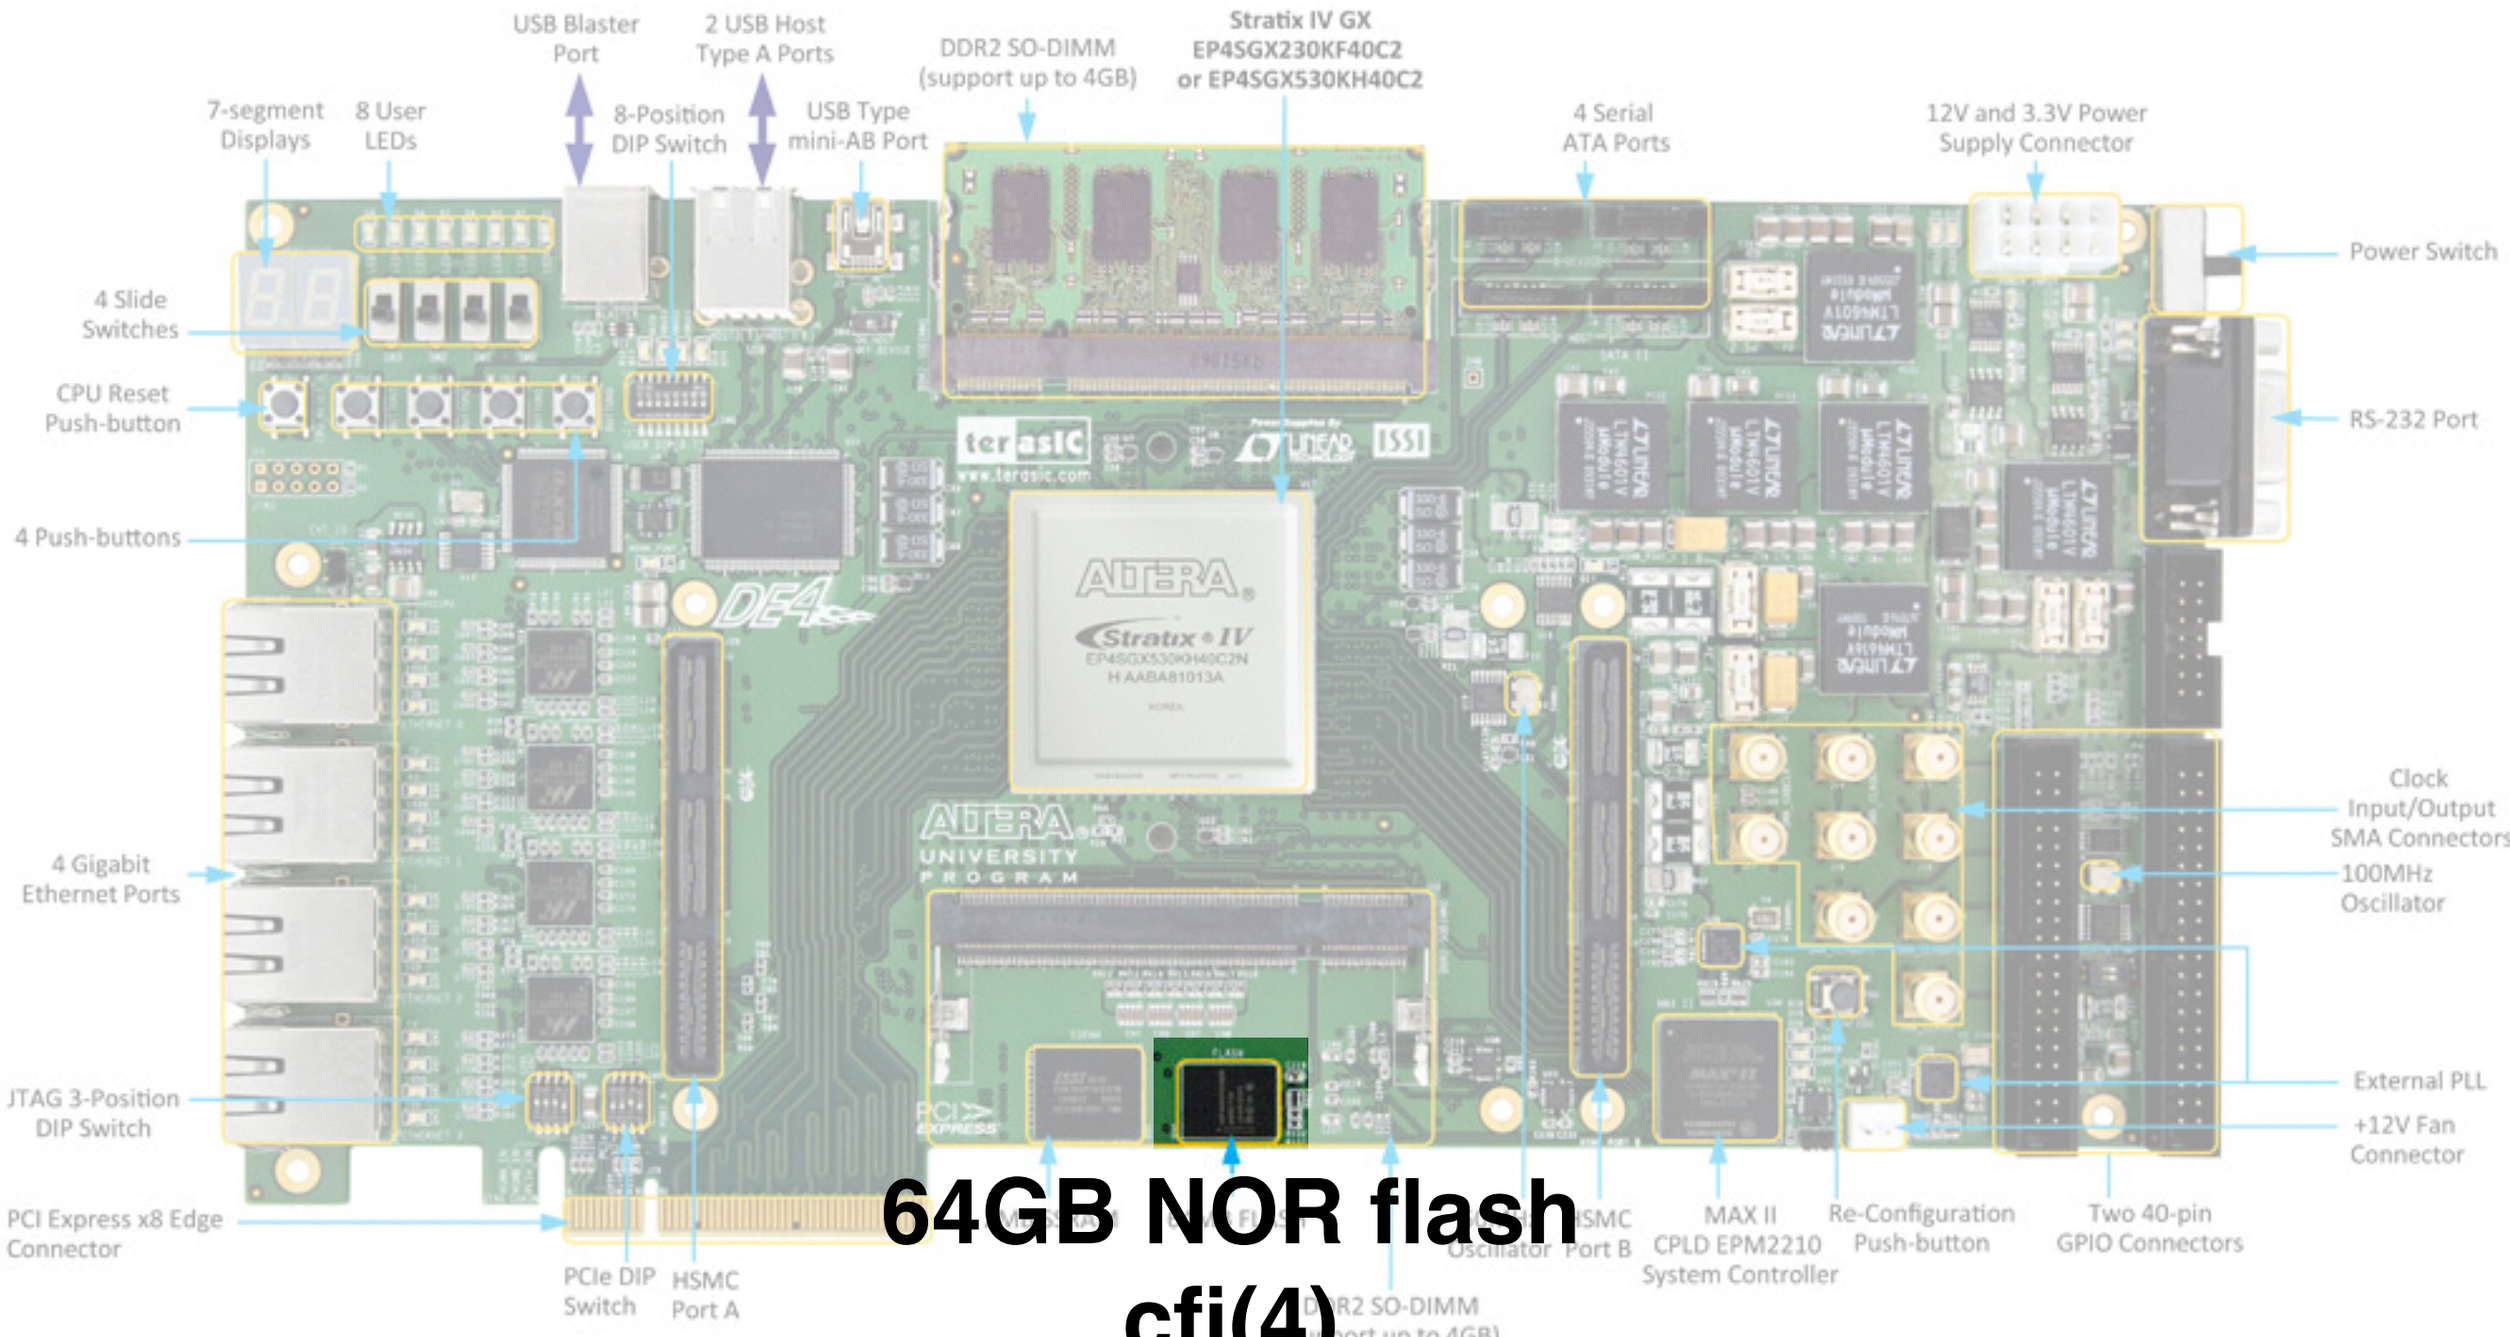
USB Blaster Port
2 USB Host Type A Ports
DDR2 SO-DIMM (support up to 4GB)
Stratix IV GX EP4SGX230KF40C2 or EP4SGX530KH40C2
7-segment Displays
8 User LEDs
8-Position DIP Switch
USB Type mini-AB Port
4 Serial ATA Ports
12V and 3.3V Power Supply Connector
Power Switch
4 Slide Switches
CPU Reset Push-button
RS-232 Port
4 Push-buttons
Clock Input/Output SMA Connectors
4 Gigabit Ethernet Ports
100MHz Oscillator
JTAG 3-Position DIP Switch
External PLL
+12V Fan Connector
PCI Express x8 Edge Connector
PCIe DIP Switch
HSMC Port A
HSMC Port B
MAX II CPLD EPM2210 System Controller
Re-Configuration Push-button
Two 40-pin GPIO Connectors
64GB NOR flash
cfi(4)

LEDs
terasic_de4led(4)

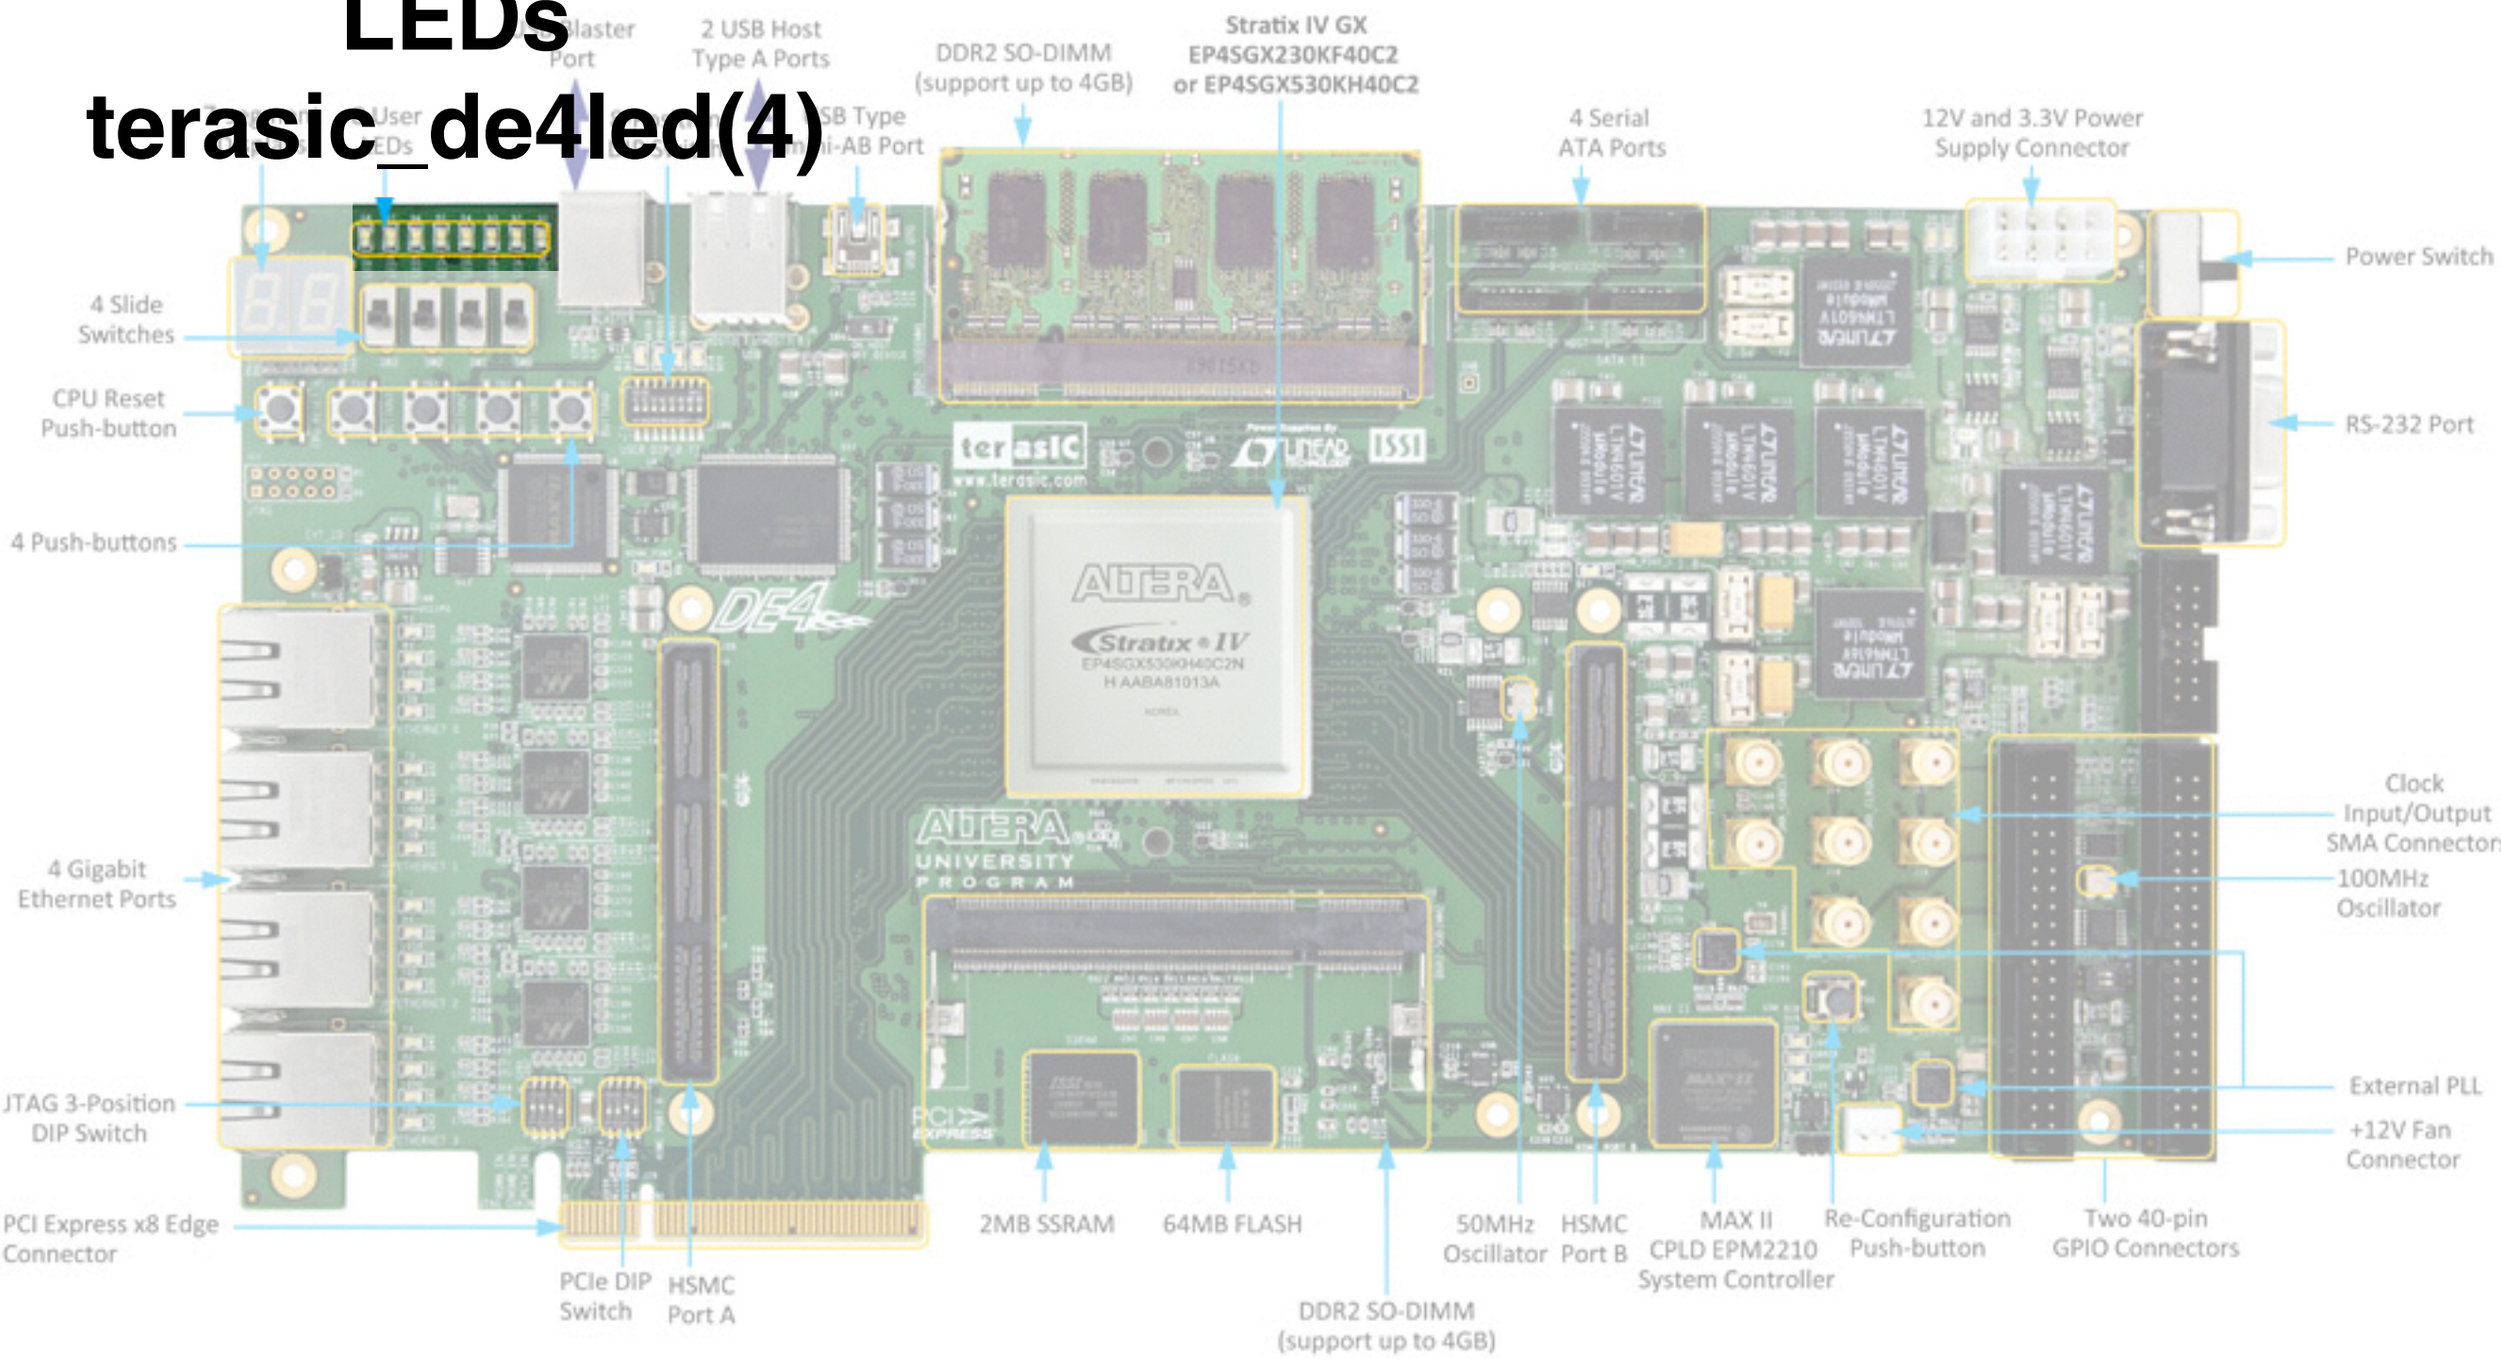

Buttons & Switches (/dev/de4bsw)

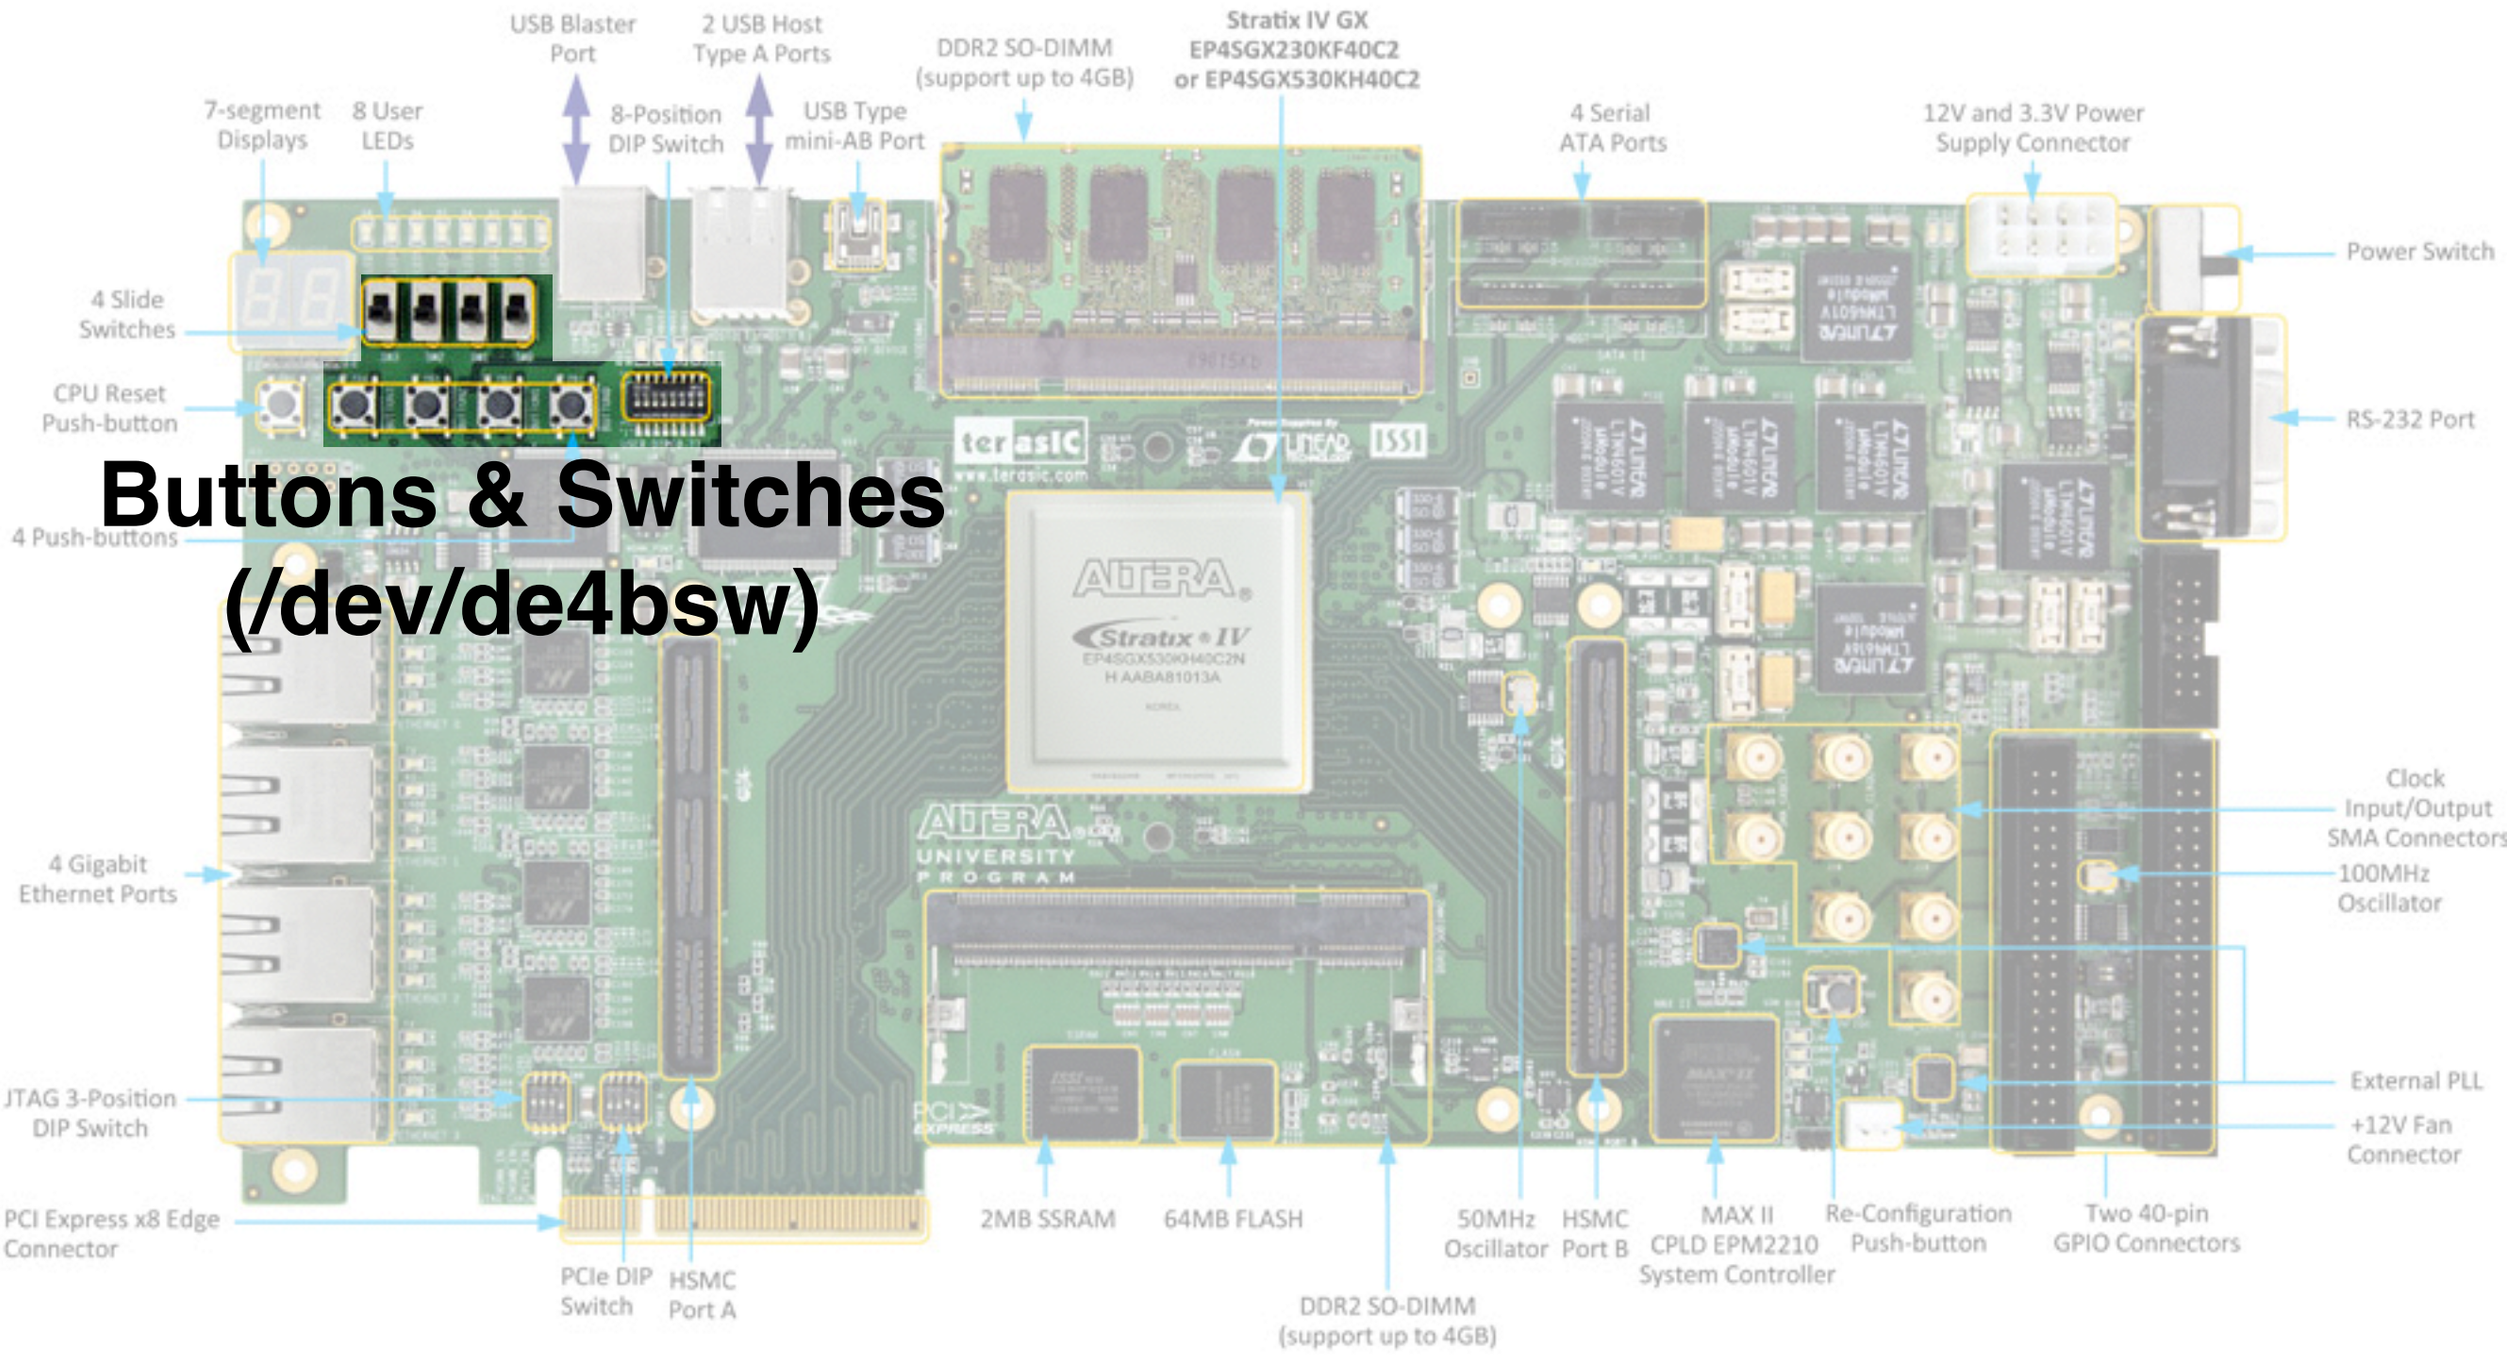

CPU temperature
/dev/de4tempfan

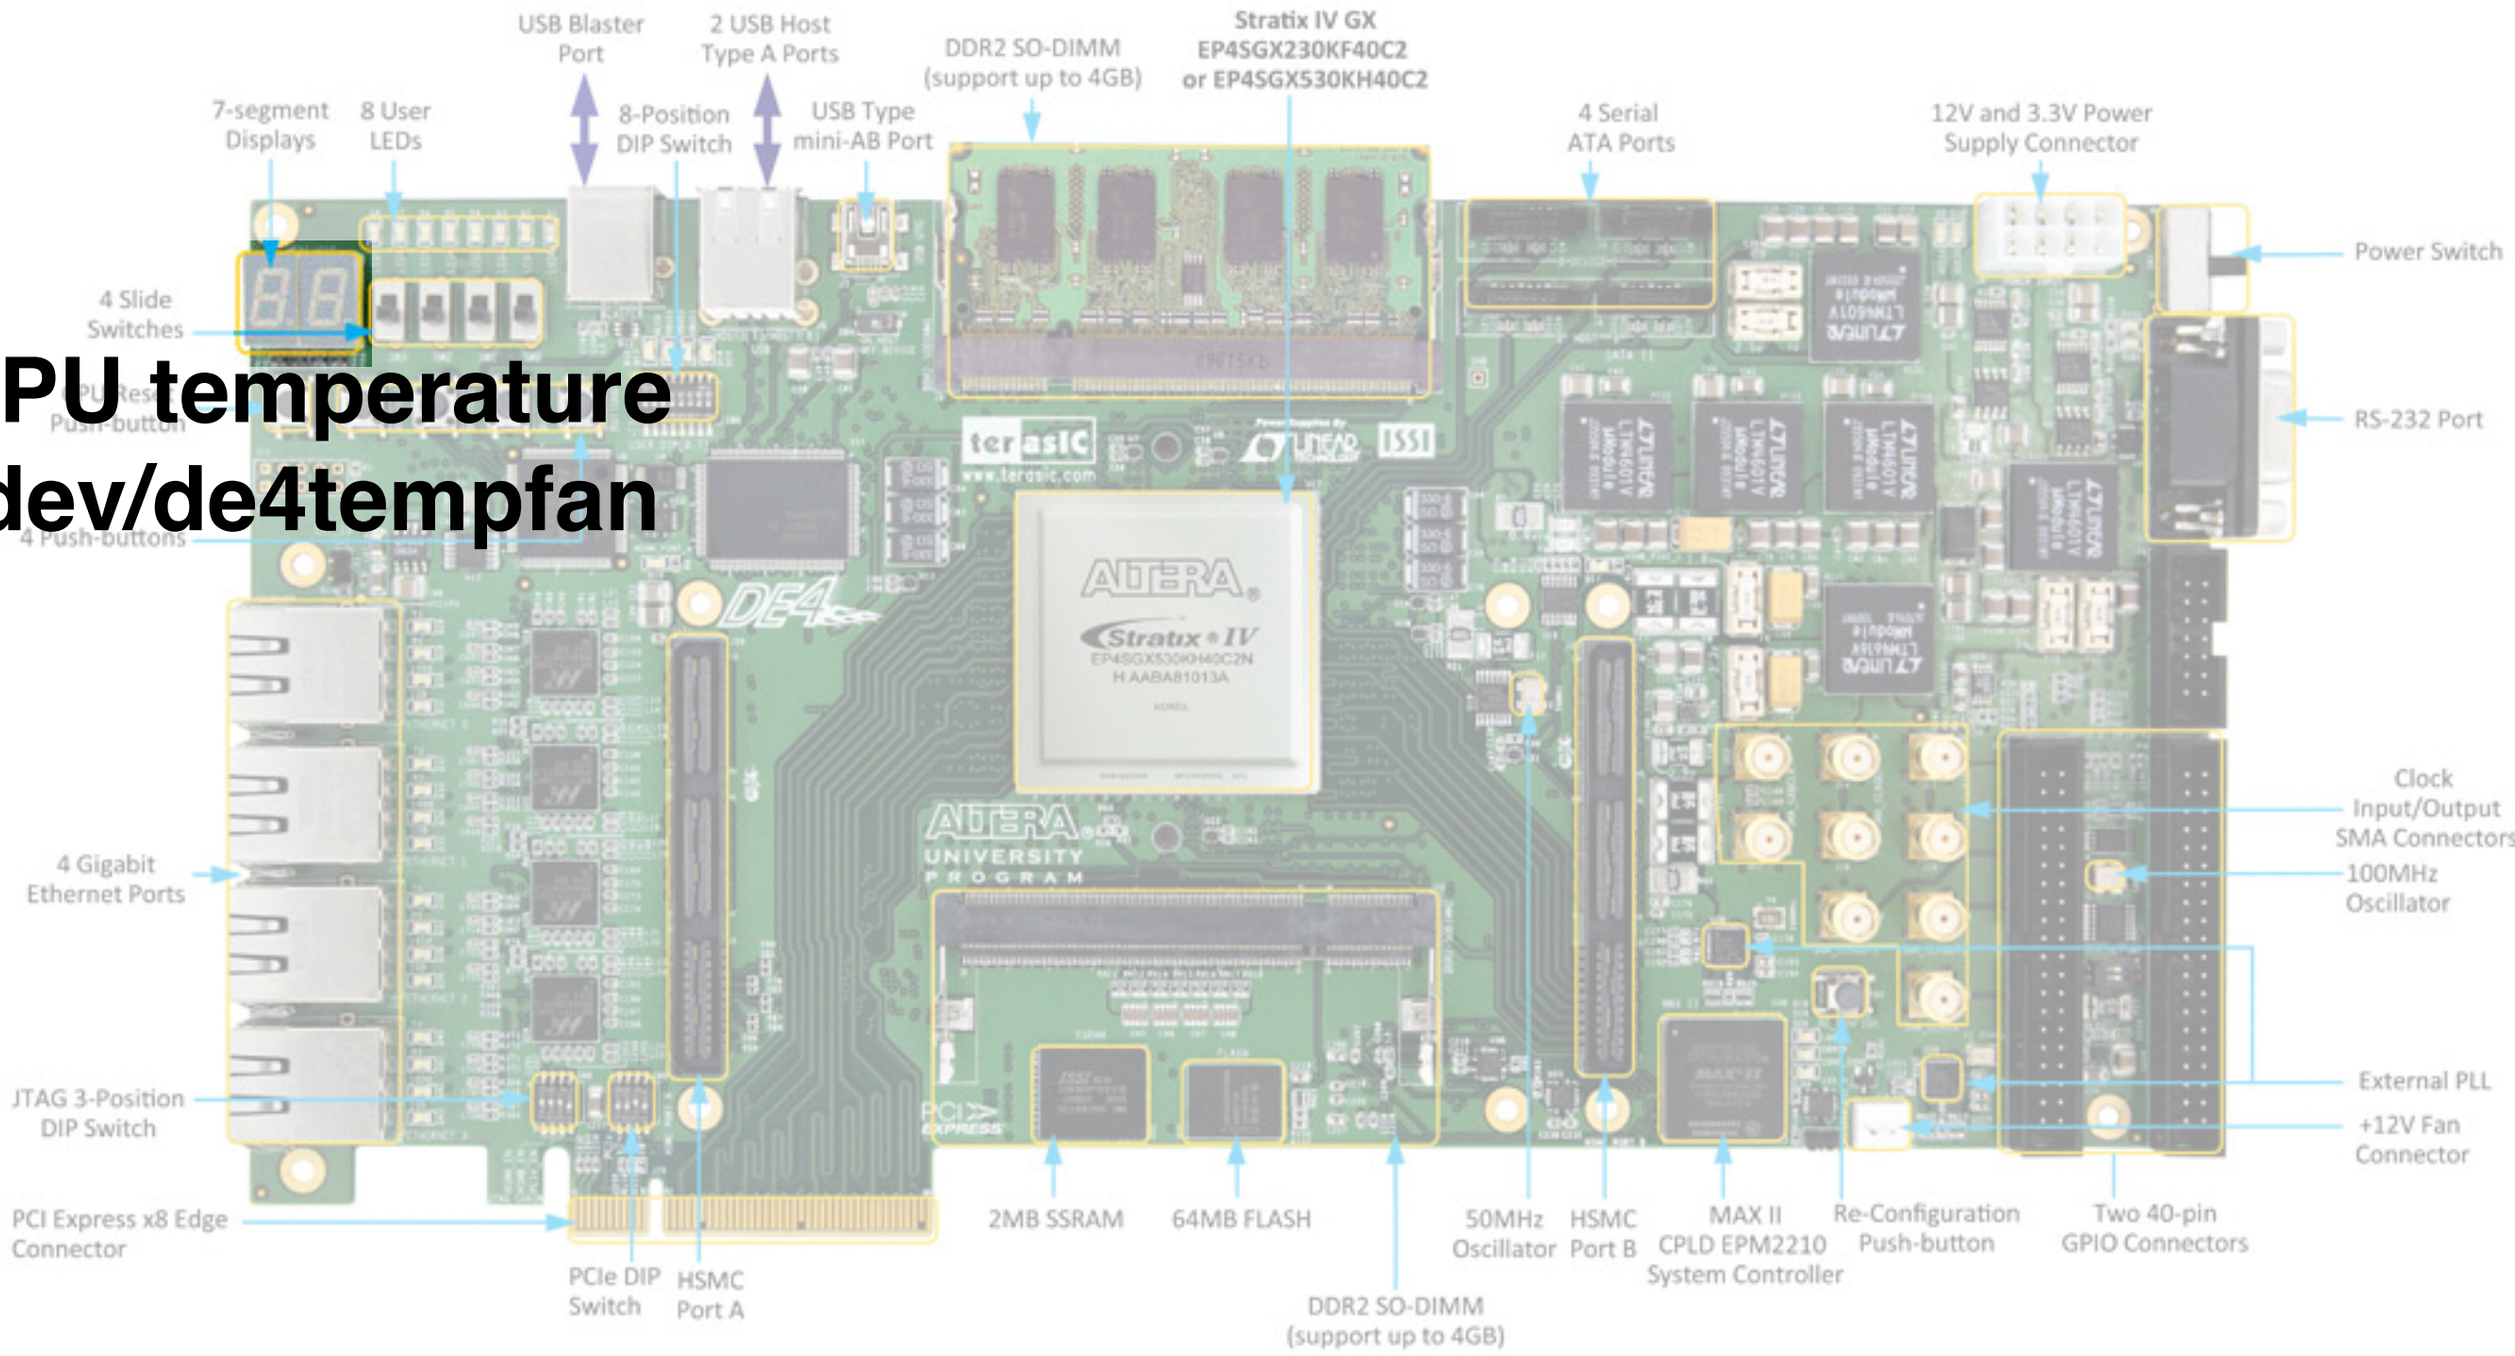

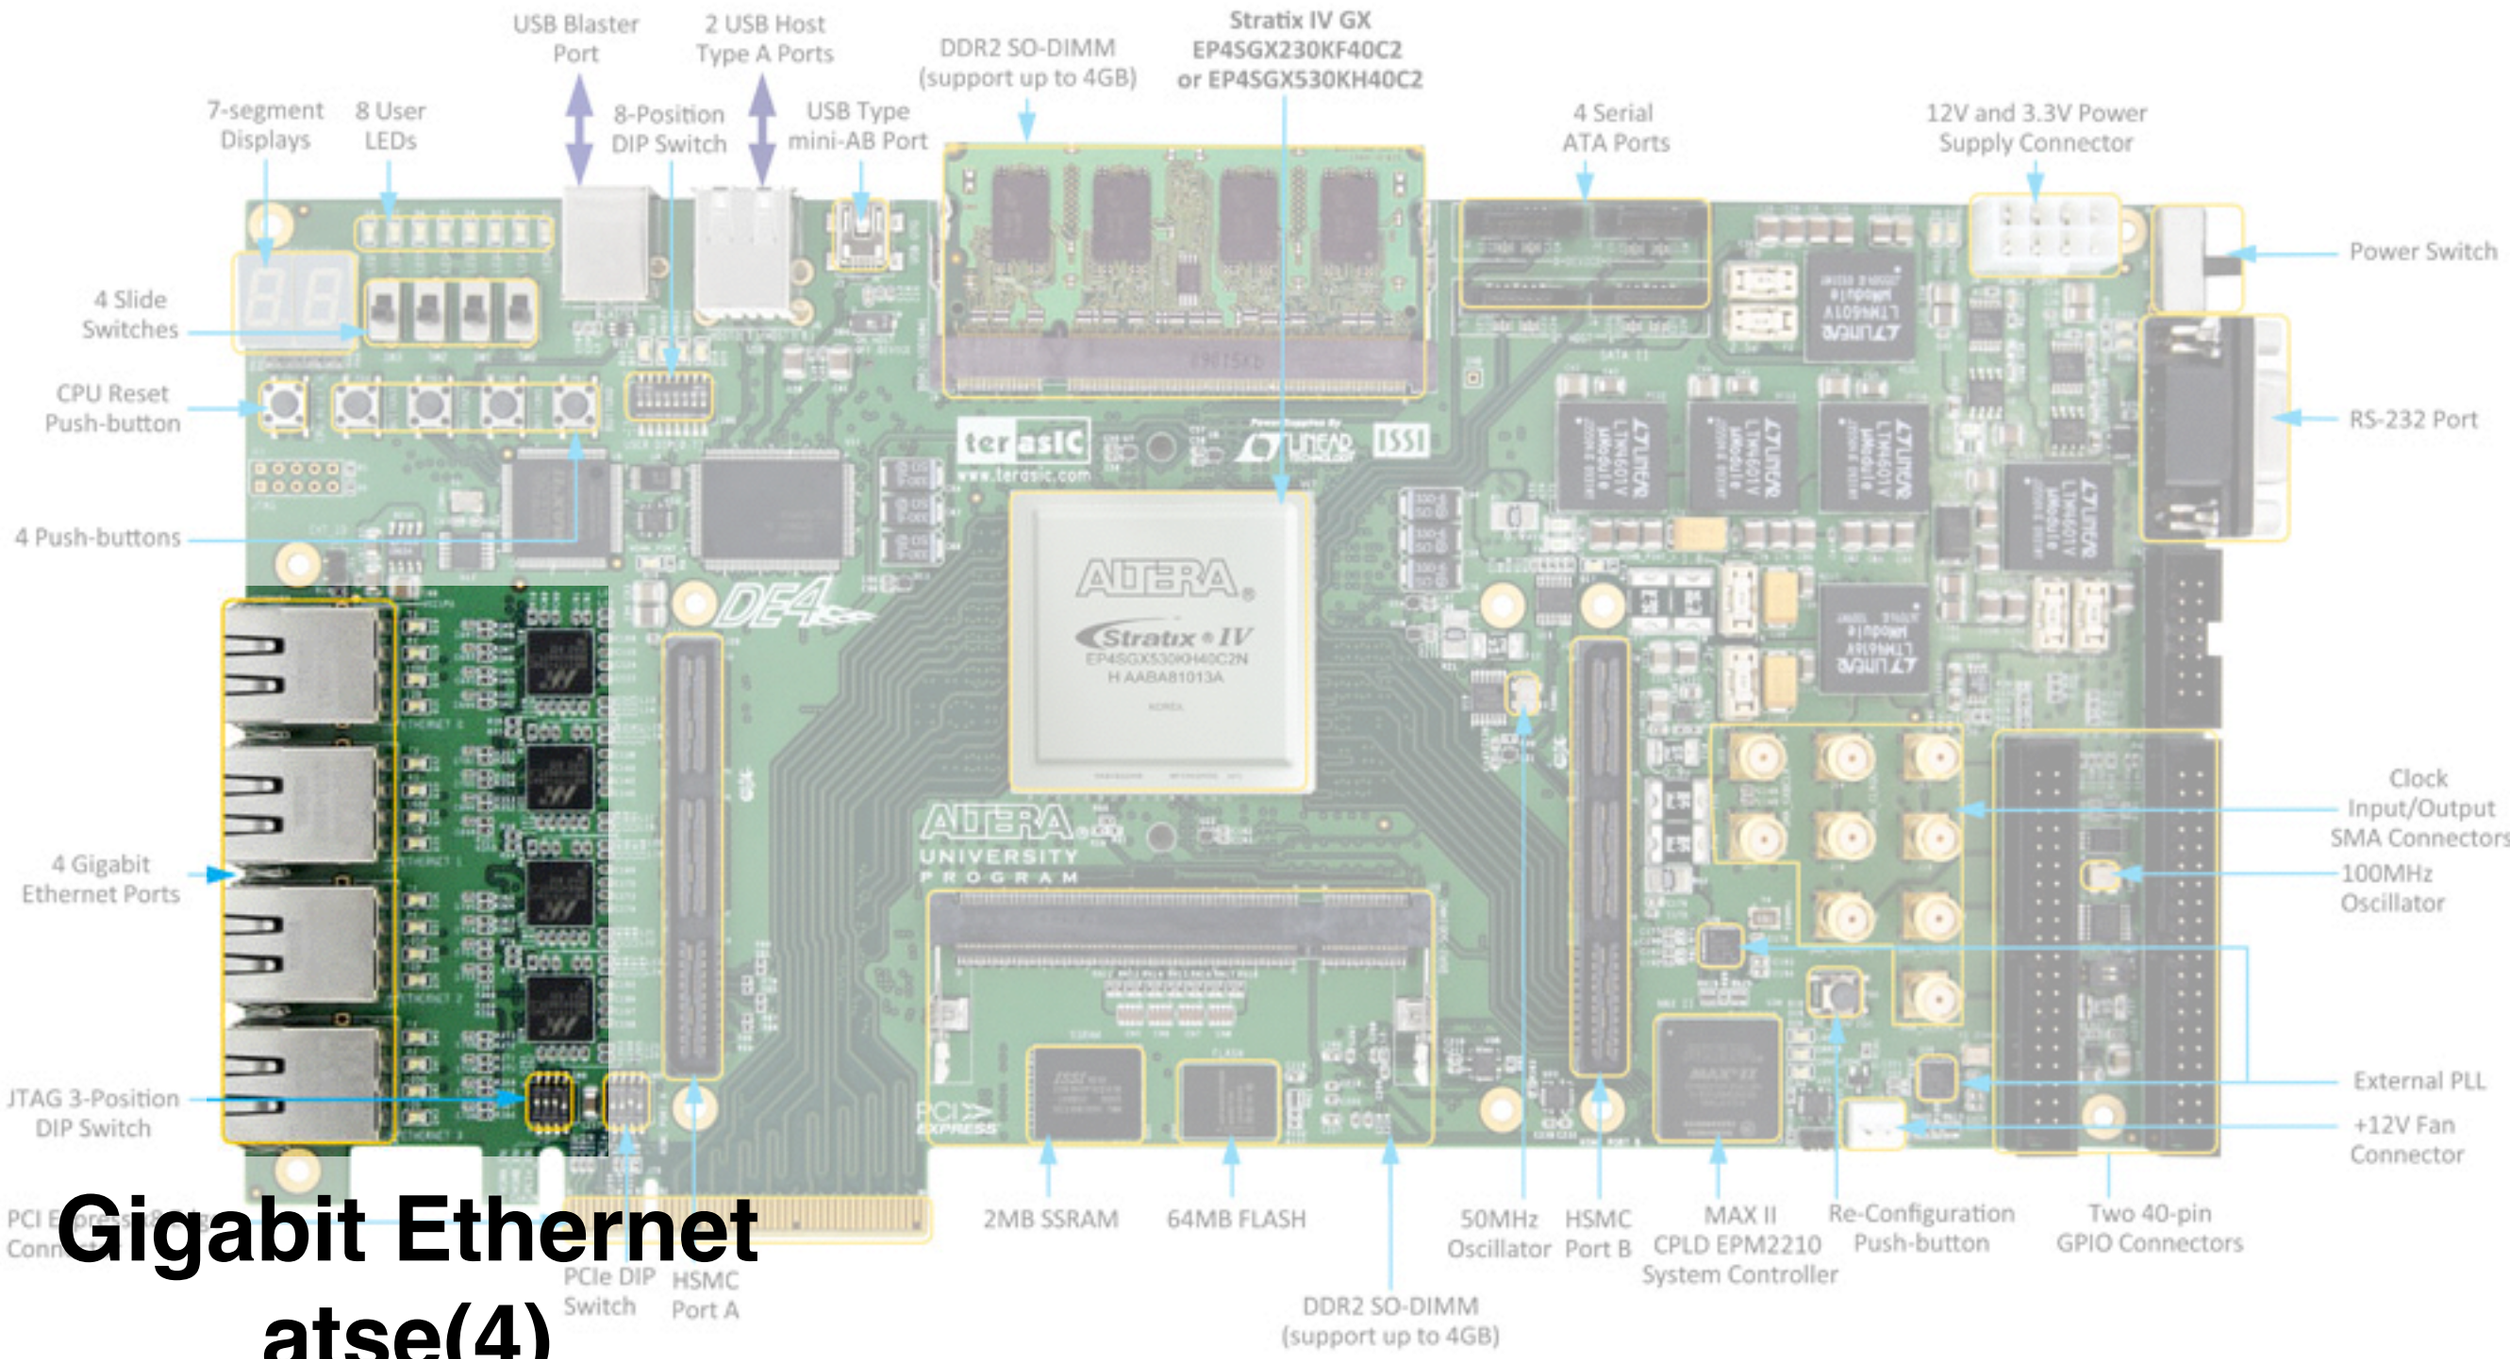

Gigabit Ethernet atse(4)

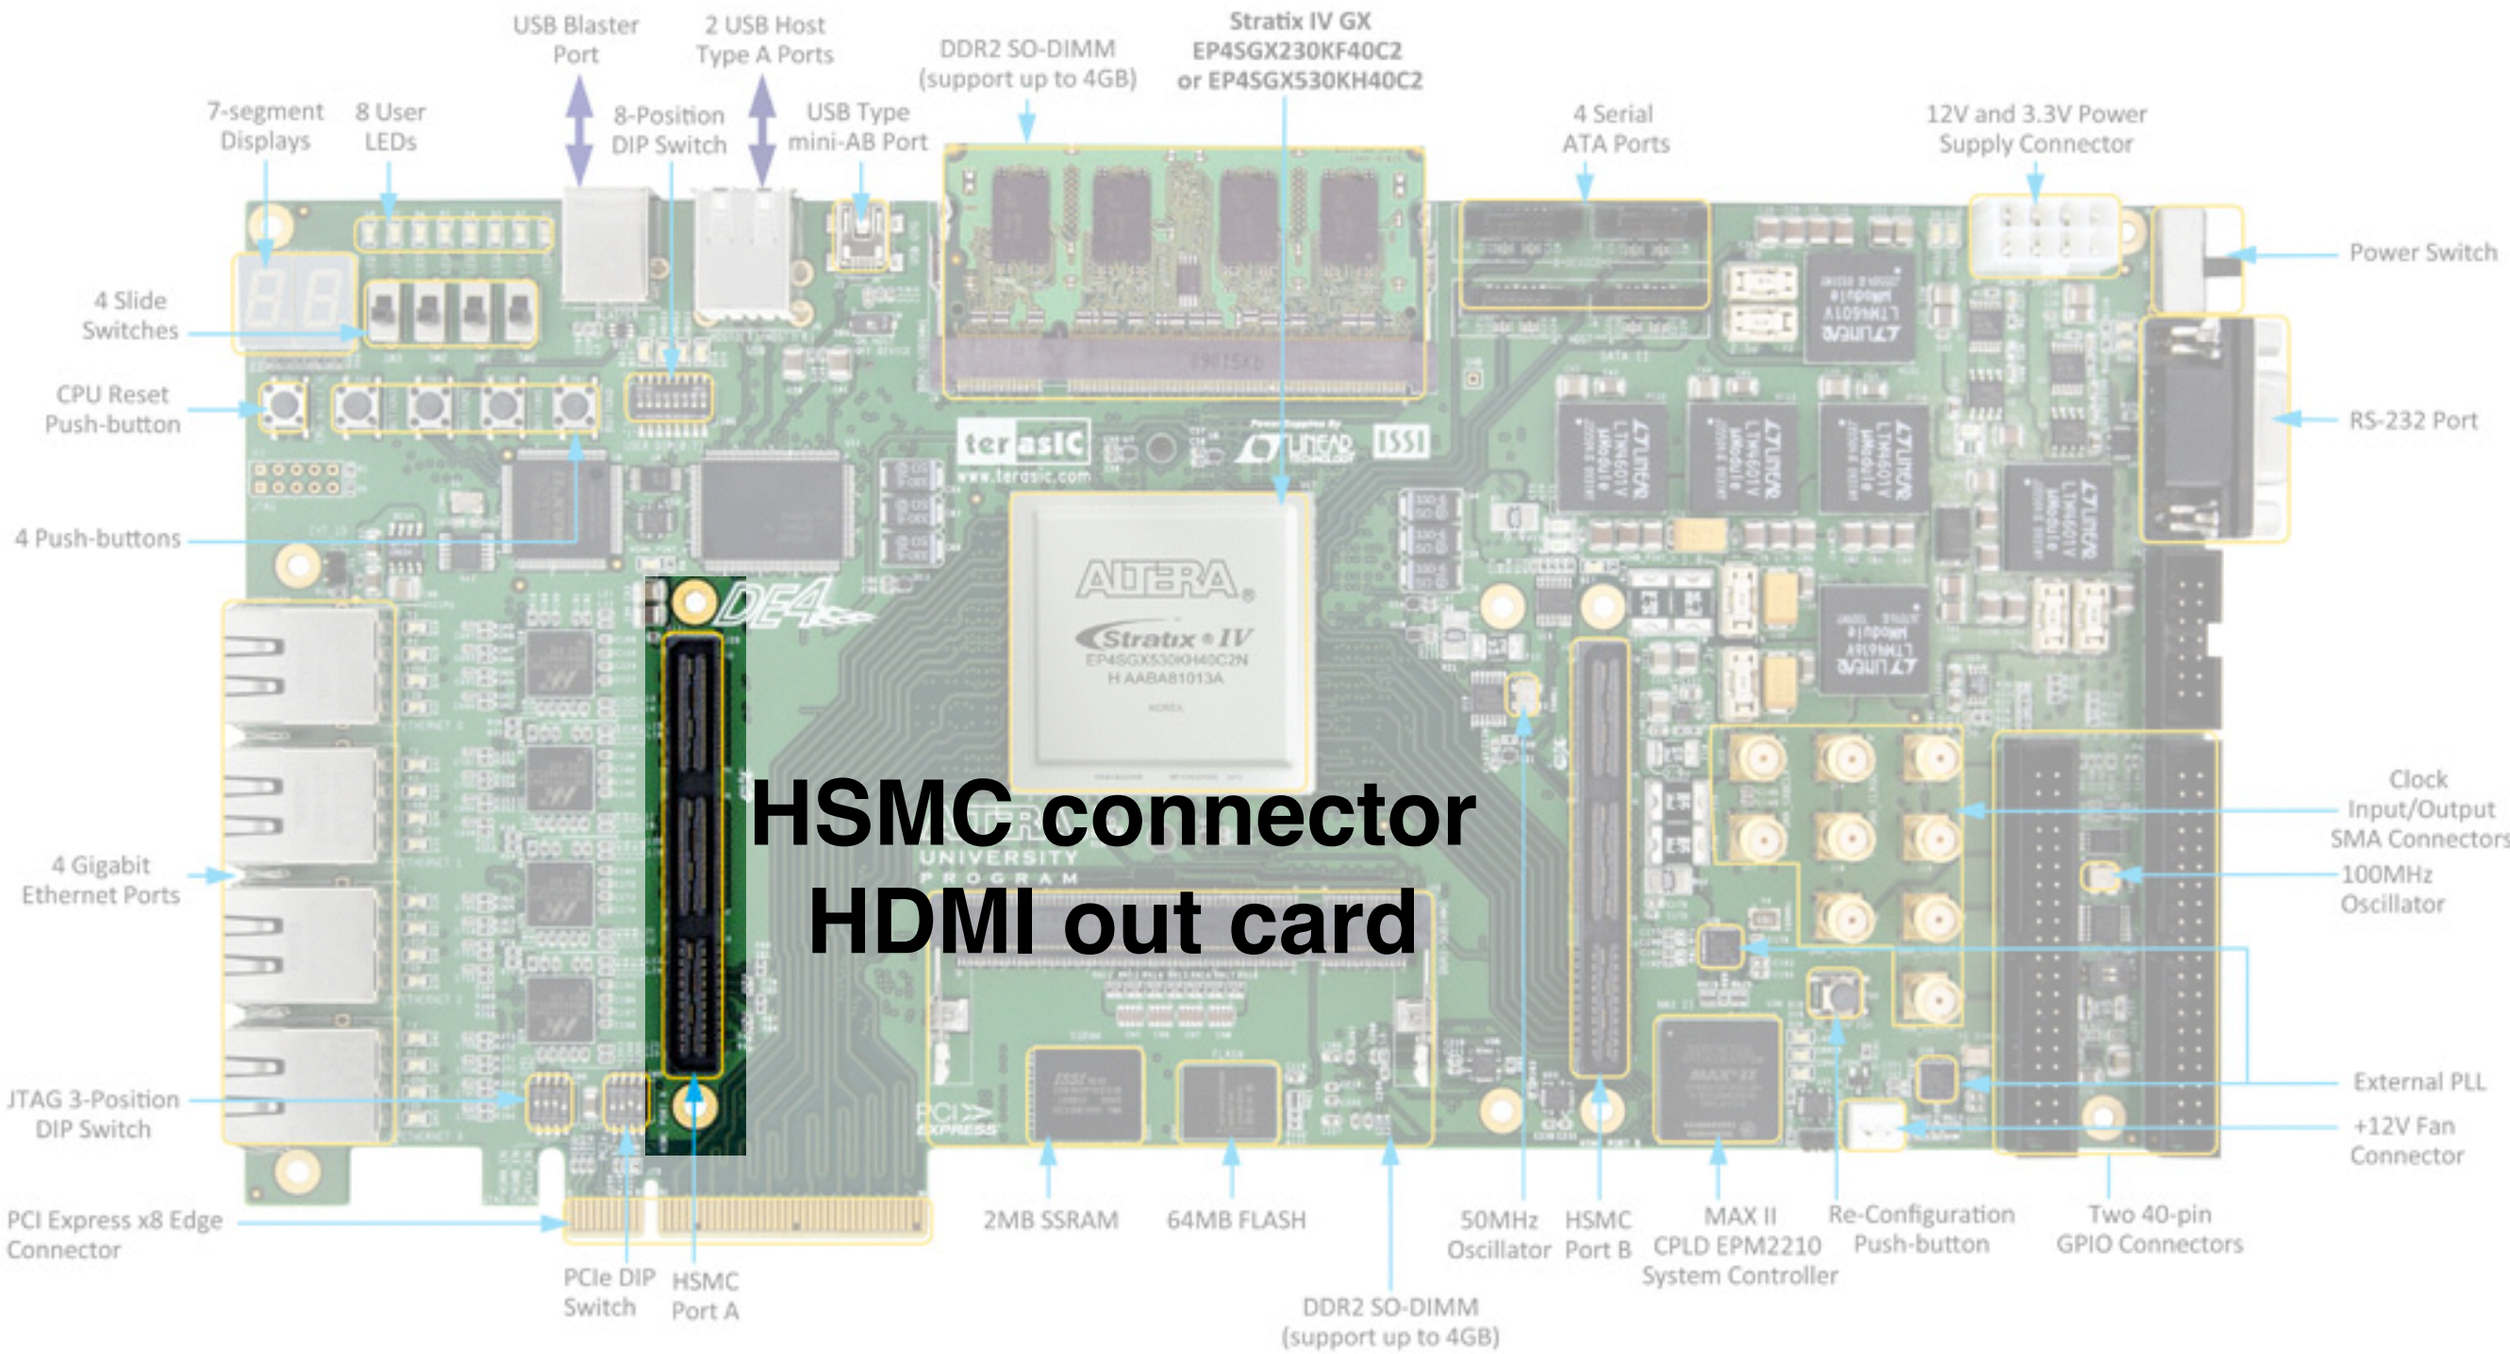
USB Blaster Port
2 USB Host Type A Ports
DDR2 SO-DIMM (support up to 4GB)
Stratix IV GX EP4SGX230KF40C2 or EP4SGX530KH40C2
7-segment Displays
8 User LEDs
8-Position DIP Switch
USB Type mini-AB Port
4 Serial ATA Ports
12V and 3.3V Power Supply Connector
Power Switch
4 Slide Switches
CPU Reset Push-button
RS-232 Port
4 Push-buttons
HSMC connector
HDMI out card
Clock Input/Output SMA Connectors
4 Gigabit Ethernet Ports
100MHz Oscillator
External PLL
JTAG 3-Position DIP Switch
+12V Fan Connector
PCI Express x8 Edge Connector
2MB SSRAM
64MB FLASH
50MHz Oscillator
HSMC Port B
MAX II CPLD EPM2210 System Controller
Re-Configuration Push-button
Two 40-pin GPIO Connectors
PCIe DIP Switch
HSMC Port A
DDR2 SO-DIMM (support up to 4GB)

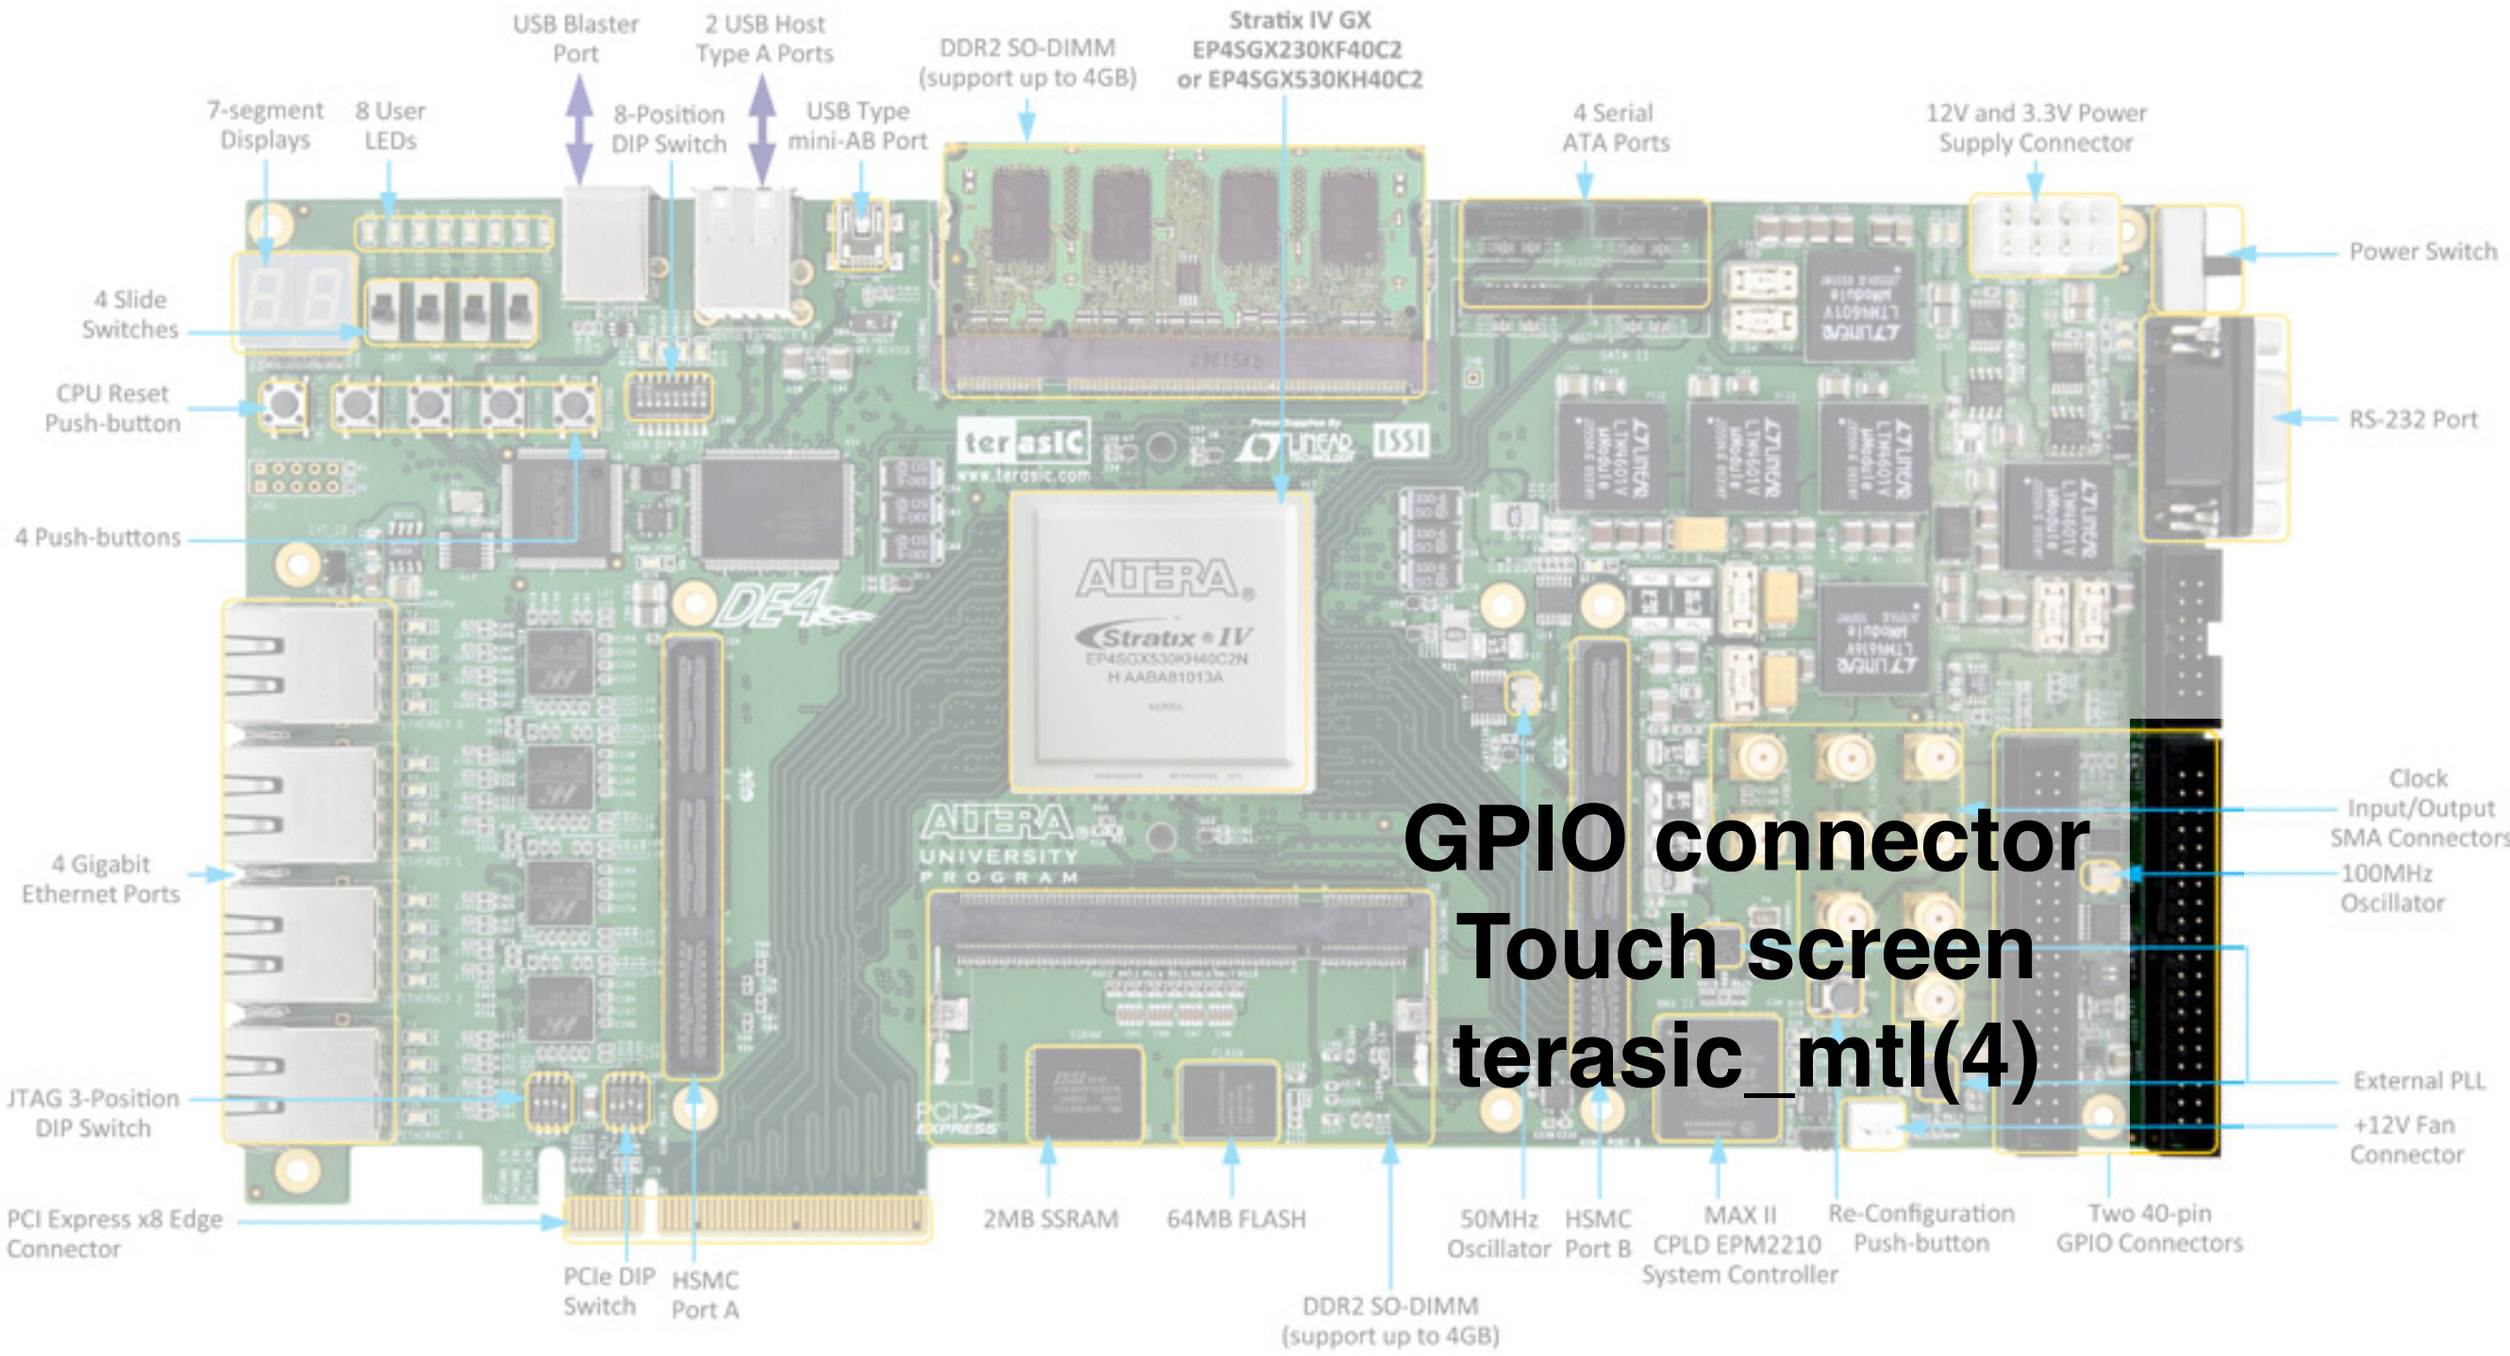
USB Blaster Port
2 USB Host Type A Ports
DDR2 SO-DIMM (support up to 4GB)
Stratix IV GX EP4SGX230KF40C2 or EP4SGX530KH40C2
7-segment Displays
8 User LEDs
8-Position DIP Switch
USB Type mini-AB Port
4 Serial ATA Ports
12V and 3.3V Power Supply Connector
Power Switch
4 Slide Switches
CPU Reset Push-button
RS-232 Port
4 Push-buttons
Clock Input/Output SMA Connectors
100MHz Oscillator
4 Gigabit Ethernet Ports
GPIO connector
Touch screen
terasic_mtl(4)
External PLL
+12V Fan Connector
JTAG 3-Position DIP Switch
PCI Express x8 Edge Connector
2MB SSRAM
64MB FLASH
50MHz Oscillator
HSMC Port B
MAX II CPLD EPM2210 System Controller
Re-Configuration Push-button
Two 40-pin GPIO Connectors
PCIe DIP Switch
HSMC Port A
DDR2 SO-DIMM (support up to 4GB)

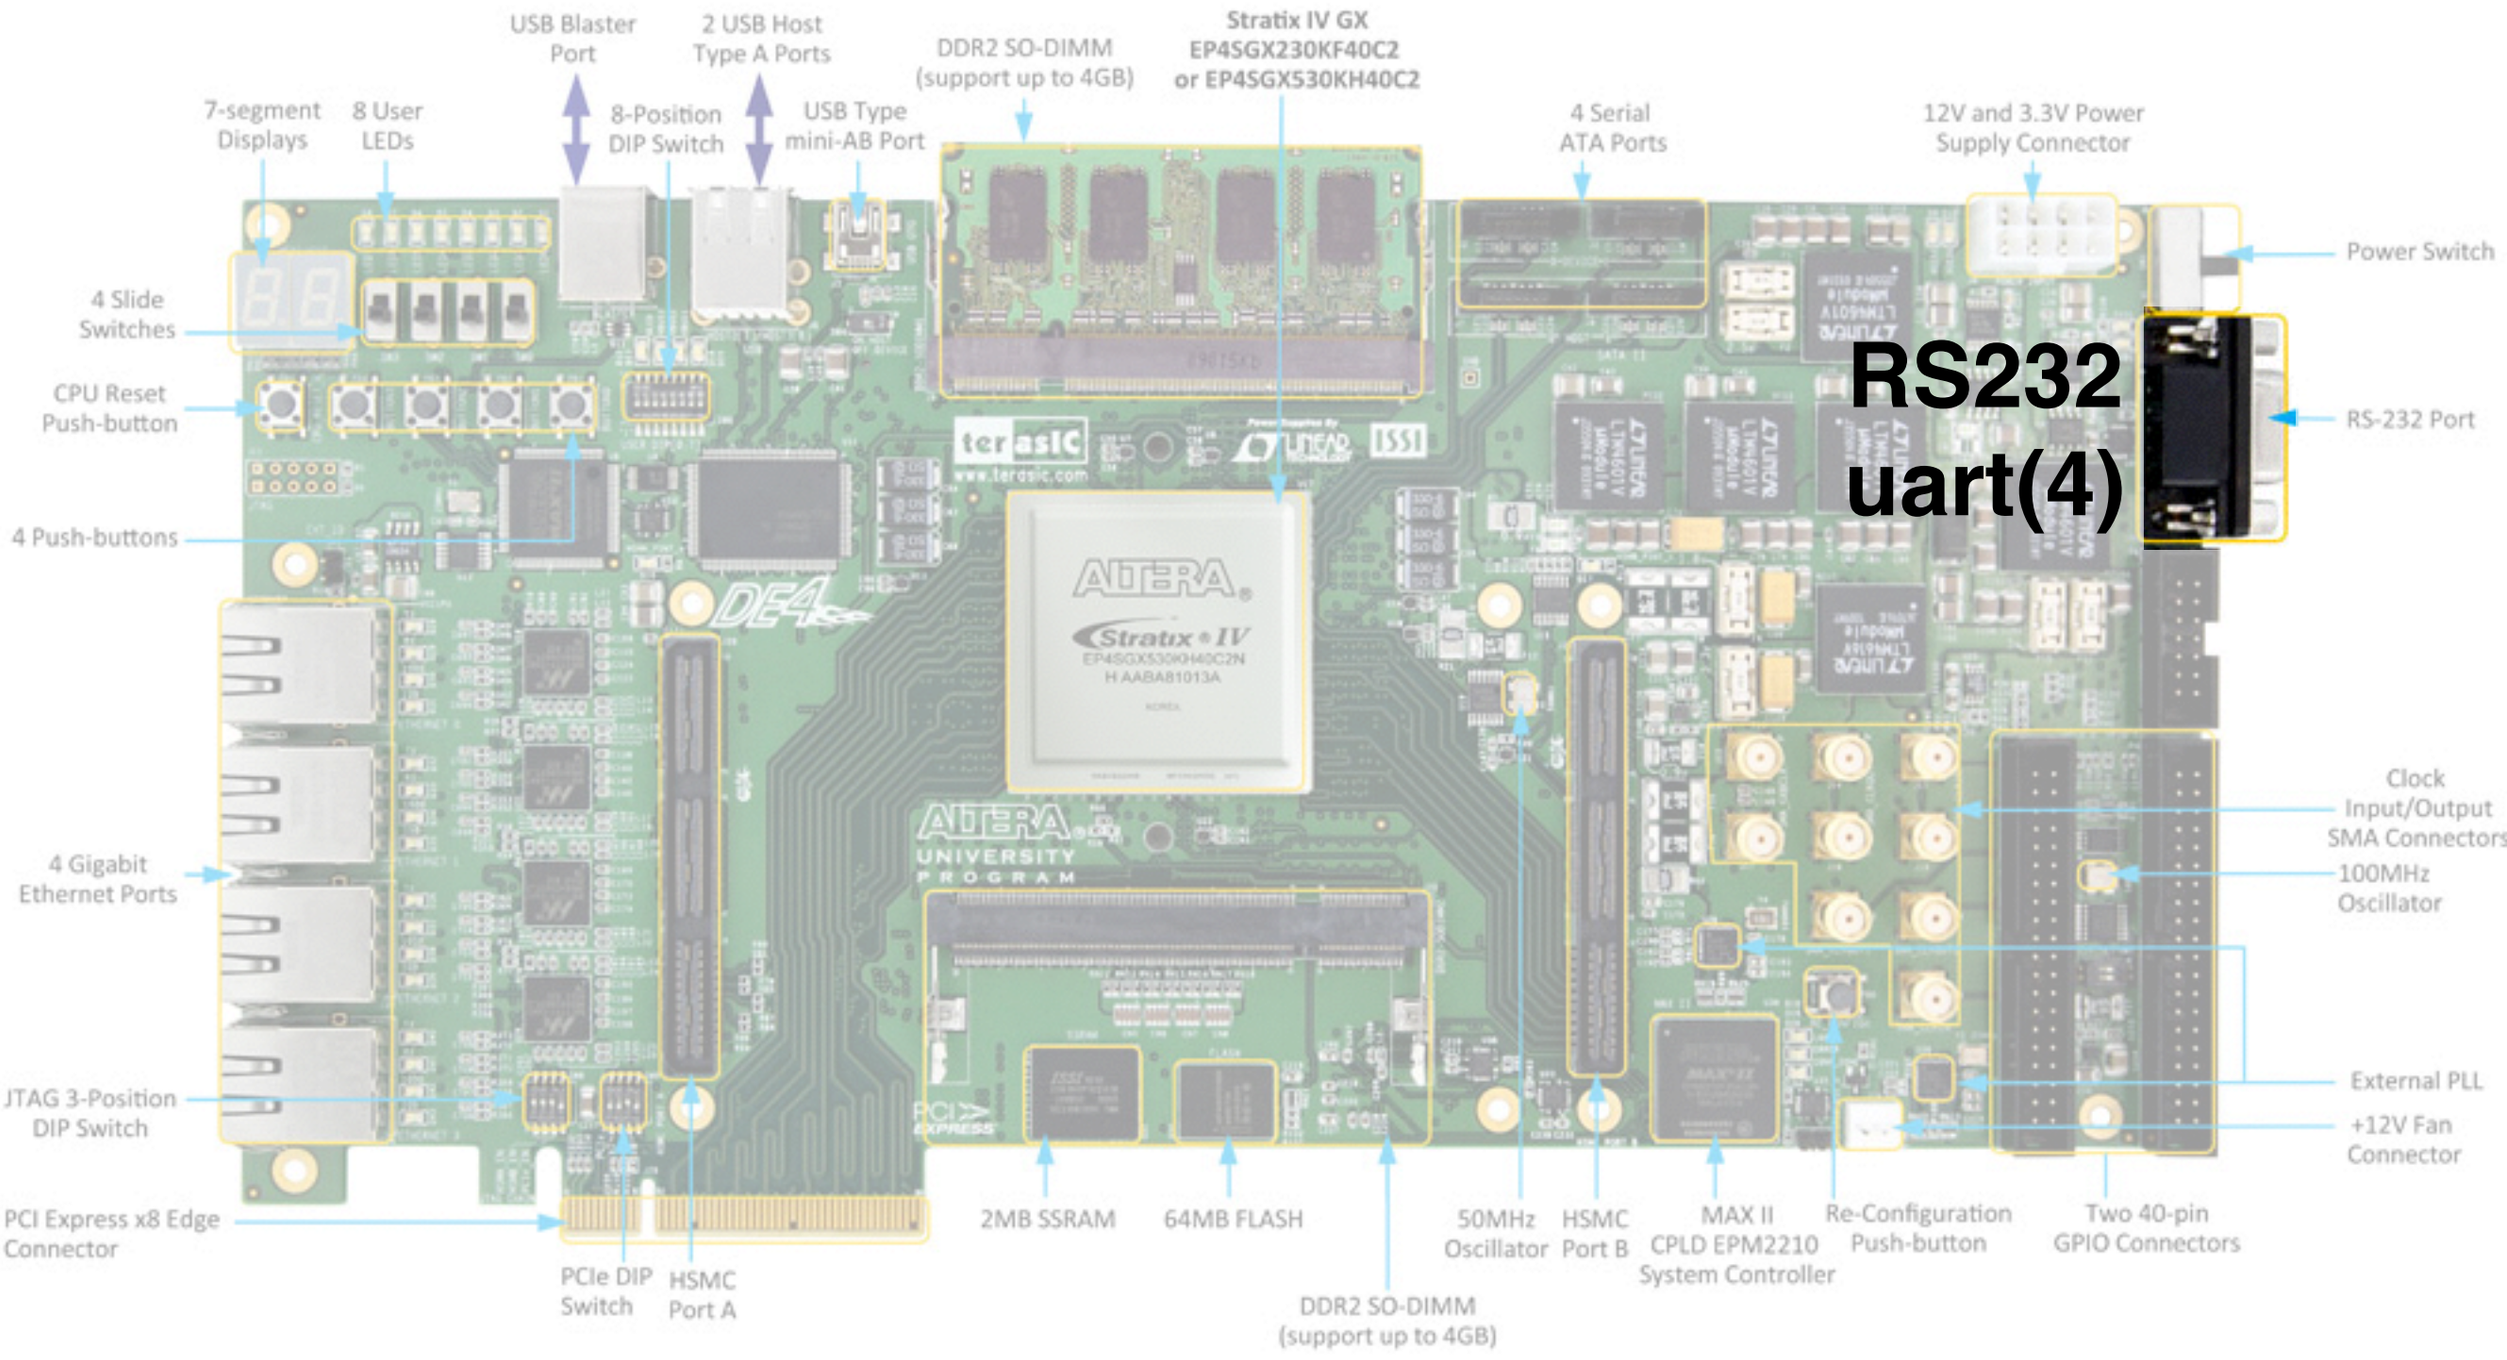
USB Blaster Port
2 USB Host Type A Ports
DDR2 SO-DIMM (support up to 4GB)
Stratix IV GX EP4SGX230KF40C2 or EP4SGX530KH40C2
7-segment Displays
8 User LEDs
8-Position DIP Switch
USB Type mini-AB Port
4 Serial ATA Ports
12V and 3.3V Power Supply Connector
Power Switch
4 Slide Switches
CPU Reset Push-button
RS232 uart(4)
RS-232 Port
4 Push-buttons
4 Gigabit Ethernet Ports
Clock Input/Output SMA Connectors
100MHz Oscillator
JTAG 3-Position DIP Switch
External PLL
+12V Fan Connector
PCI Express x8 Edge Connector
2MB SSRAM
64MB FLASH
50MHz Oscillator
HSMC Port B
MAX II CPLD EPM2210 System Controller
Re-Configuration Push-button
Two 40-pin GPIO Connectors
PCIe DIP Switch
HSMC Port A
DDR2 SO-DIMM (support up to 4GB)

SD card (2GB)
altera_sdcard(4)

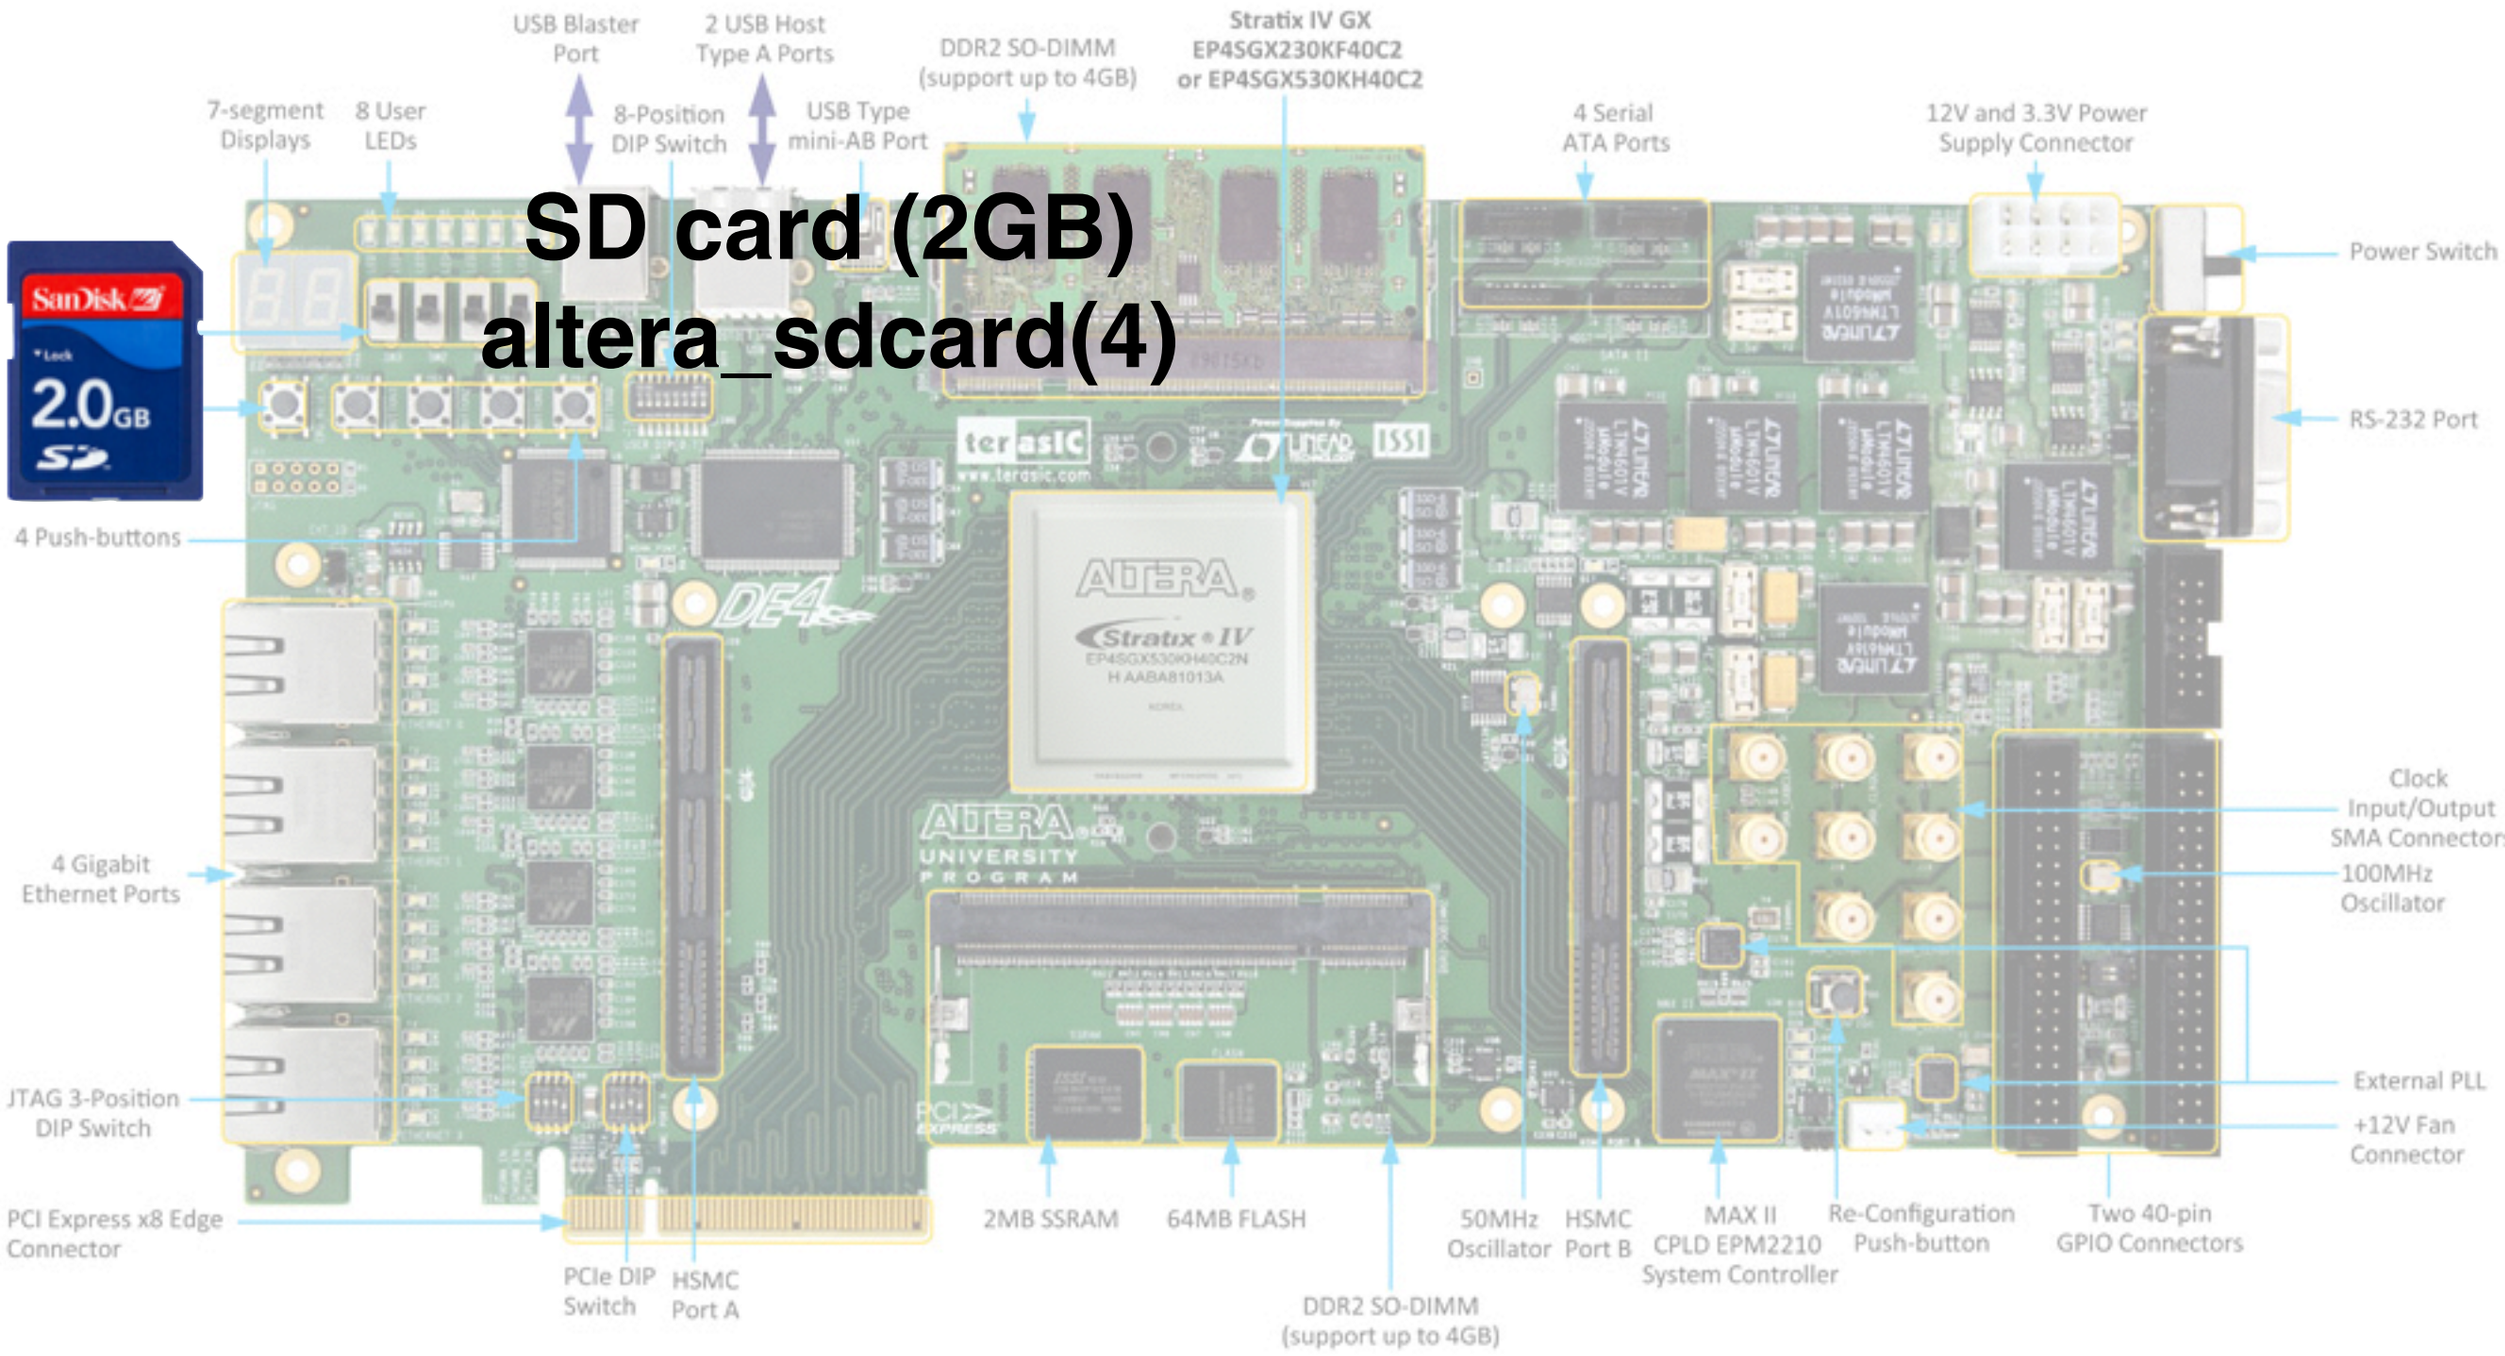

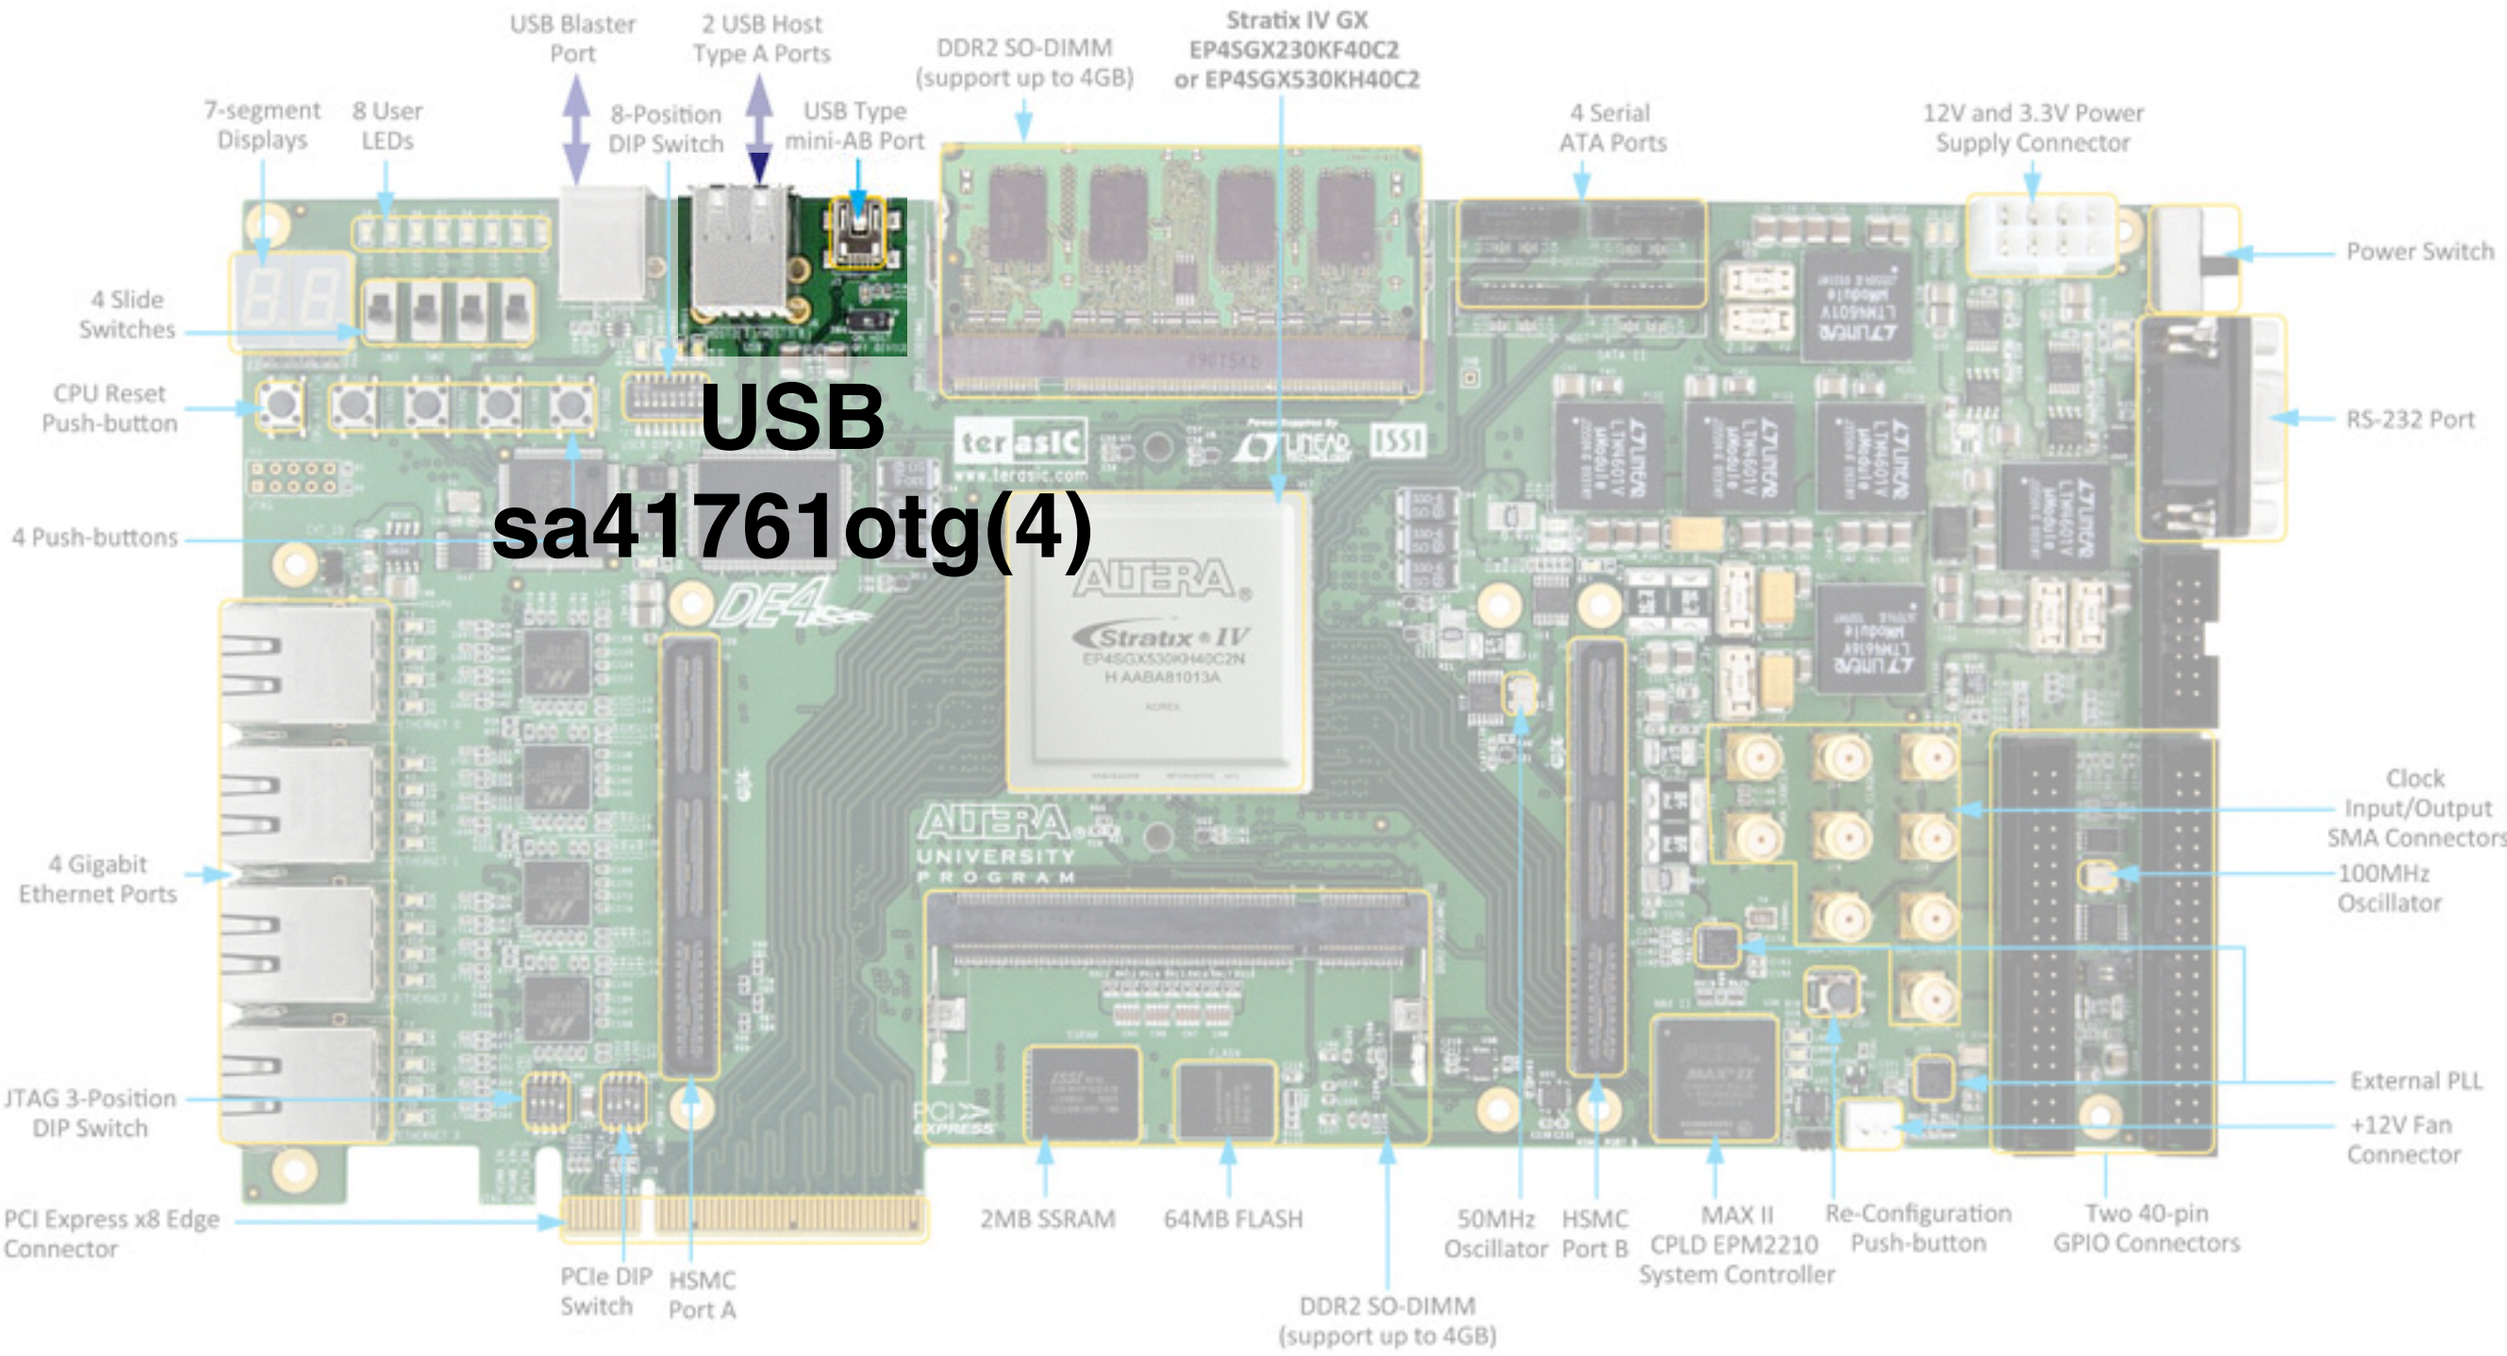
USB Blaster Port
2 USB Host Type A Ports
DDR2 SO-DIMM (support up to 4GB)
Stratix IV GX EP4SGX230KF40C2 or EP4SGX530KH40C2
7-segment Displays
8 User LEDs
8-Position DIP Switch
USB Type mini-AB Port
4 Serial ATA Ports
12V and 3.3V Power Supply Connector
Power Switch
4 Slide Switches
CPU Reset Push-button
RS-232 Port
USB sa41761otg(4)
4 Push-buttons
4 Gigabit Ethernet Ports
Clock Input/Output SMA Connectors
100MHz Oscillator
JTAG 3-Position DIP Switch
External PLL
+12V Fan Connector
PCI Express x8 Edge Connector
2MB SSRAM
64MB FLASH
50MHz Oscillator
HSMC Port B
MAX II CPLD EPM2210 System Controller
Re-Configuration Push-button
Two 40-pin GPIO Connectors
PCIe DIP Switch
HSMC Port A
DDR2 SO-DIMM (support up to 4GB)


Booting CheriBSD
Please Wait
UNIVERSITY OF CAMBRIDGE
ON OFF
UNIVERSITY OF CAMBRIDGE

Booting the DE4

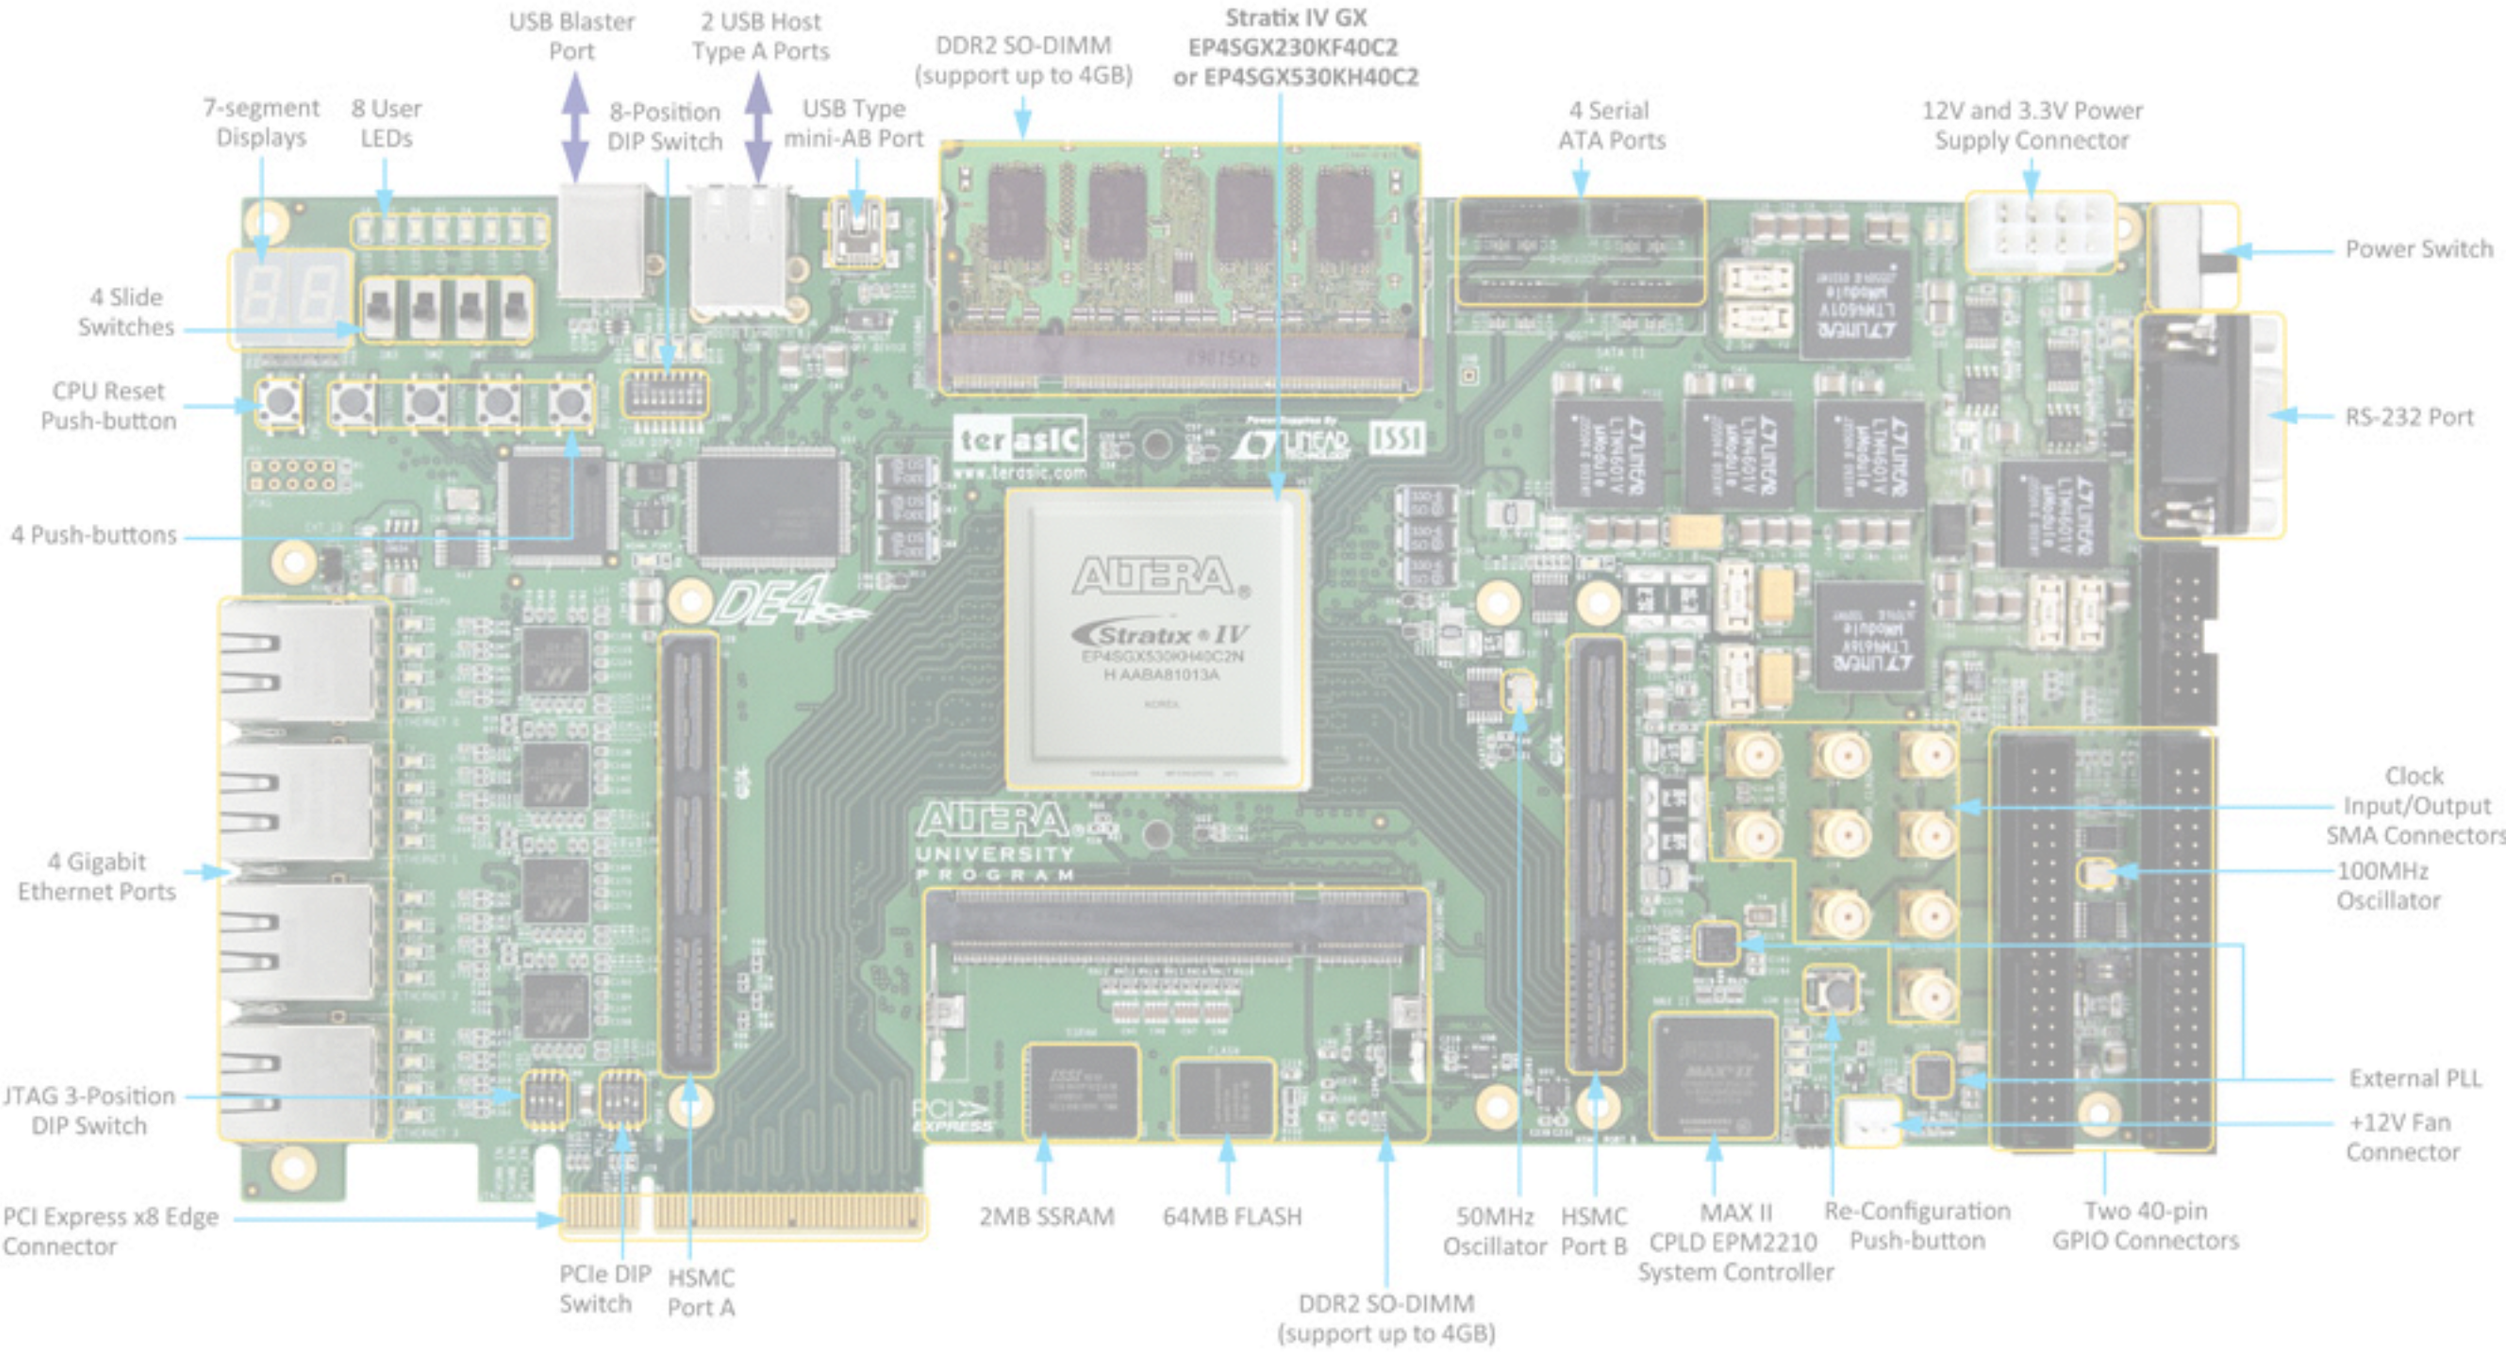

Booting the DE4

DRAM

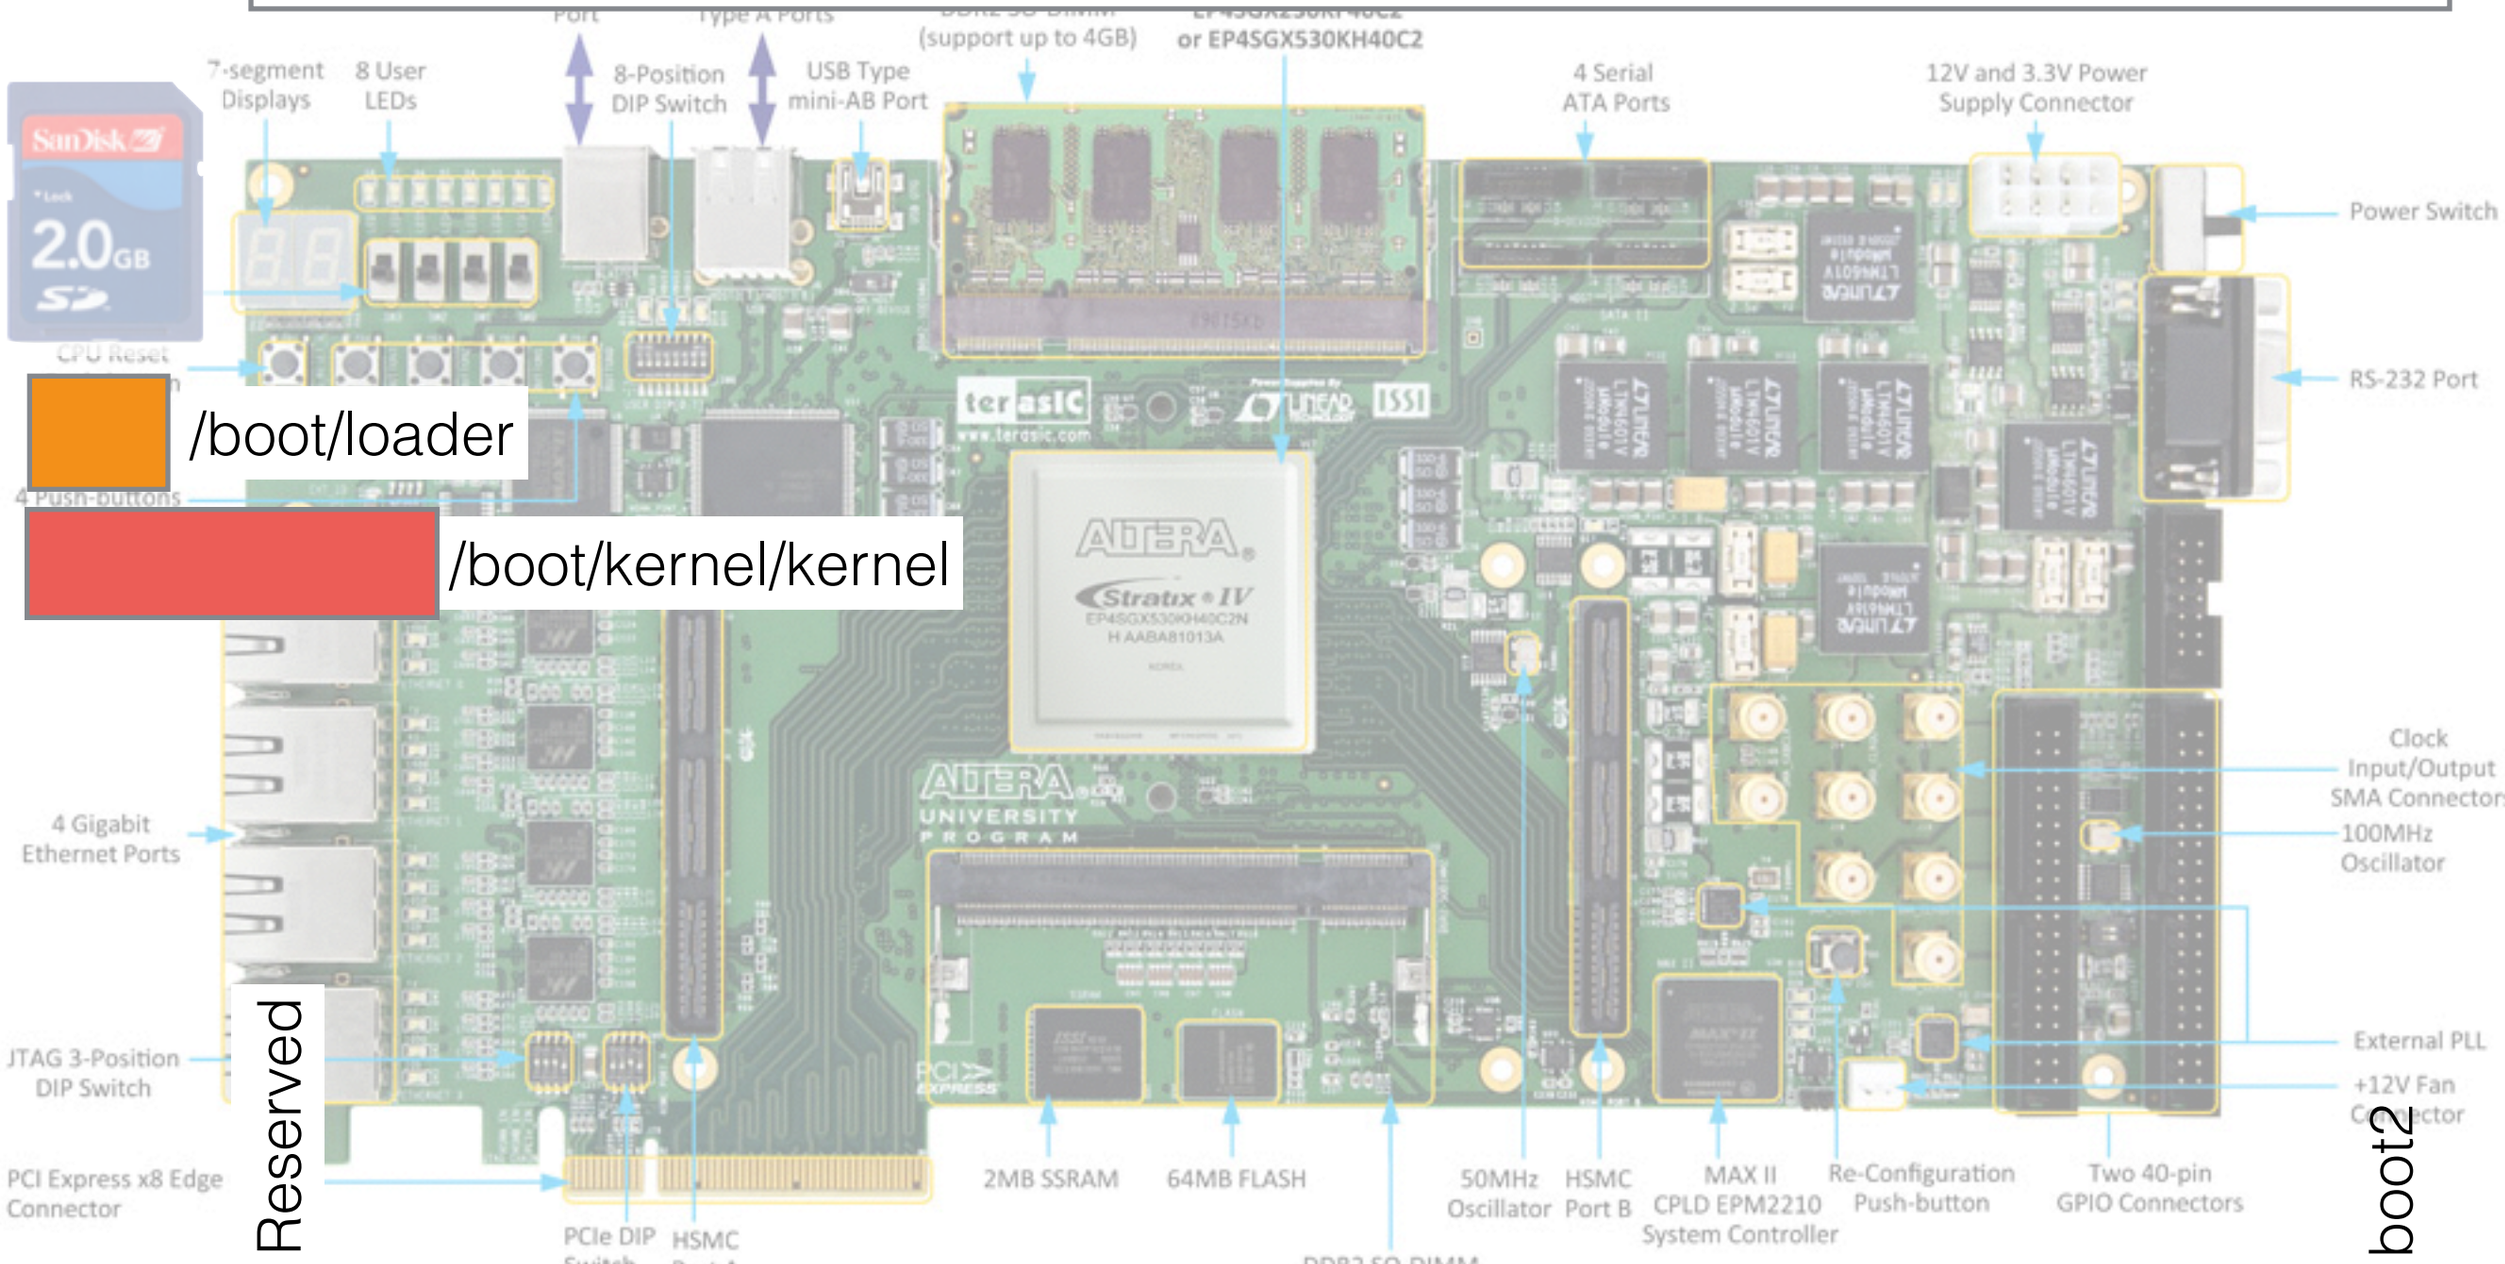

Flash

	FPGA 1	FPGA2	Operating system (UFS)	
128k	12MB	12MB	39.75MB	128k

Booting the DE4

DRAM

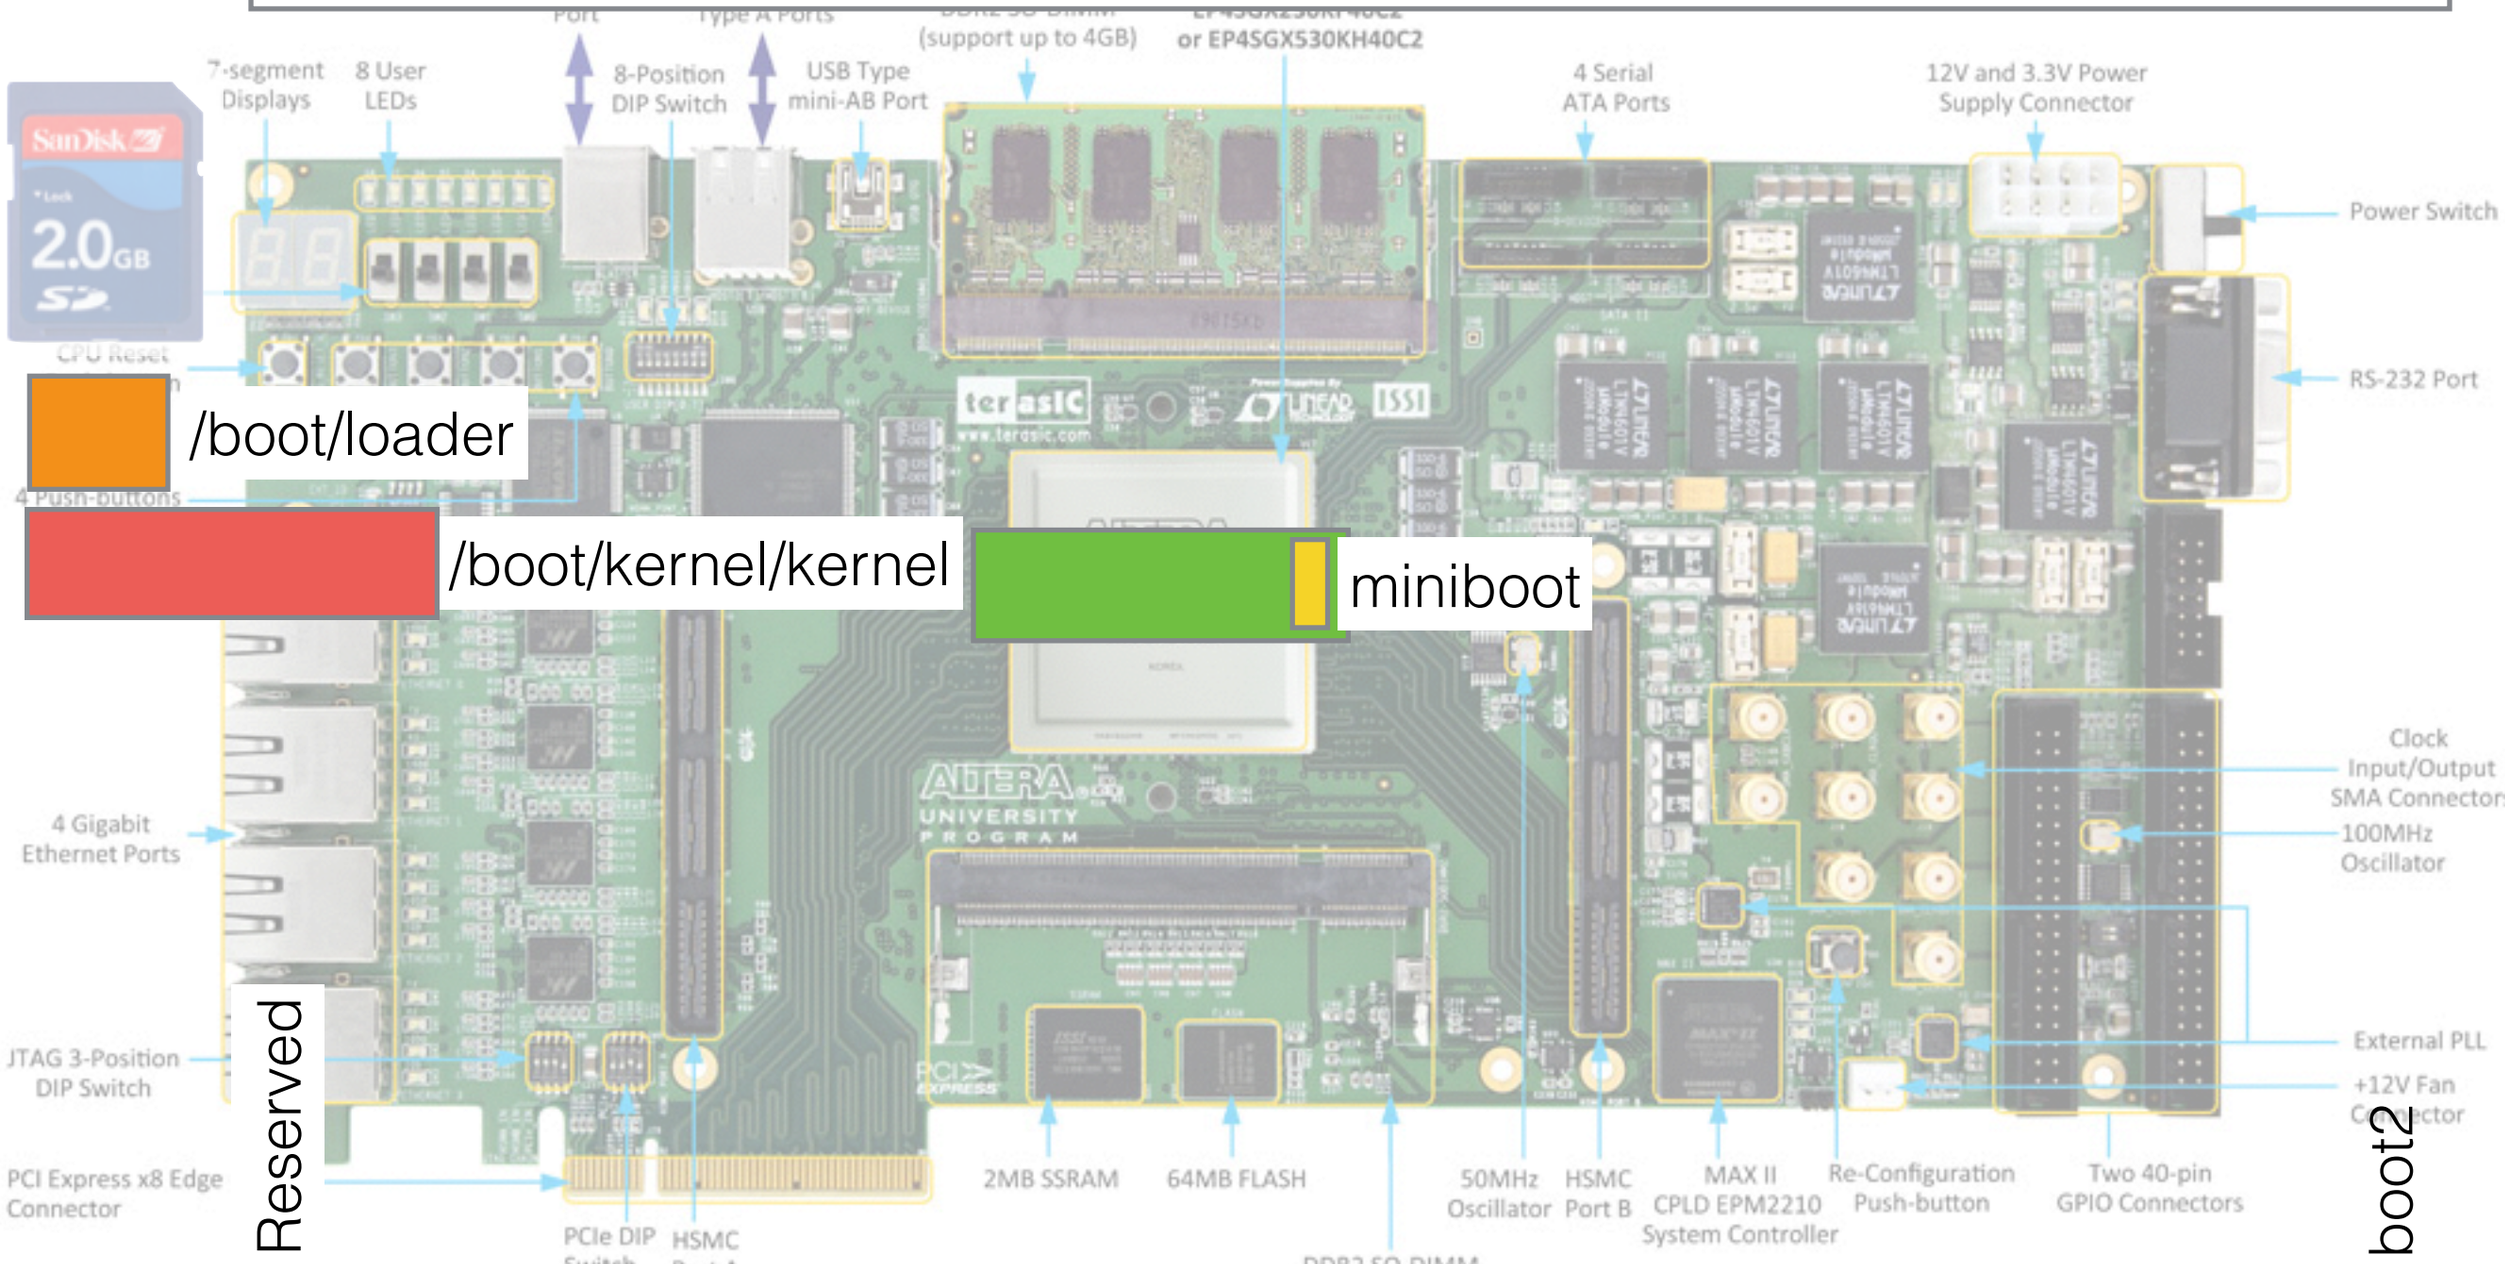

Flash

	FPGA 1	FPGA2	Operating system (UFS)	
128k	12MB	12MB	39.75MB	128k

Booting the DE4

DRAM

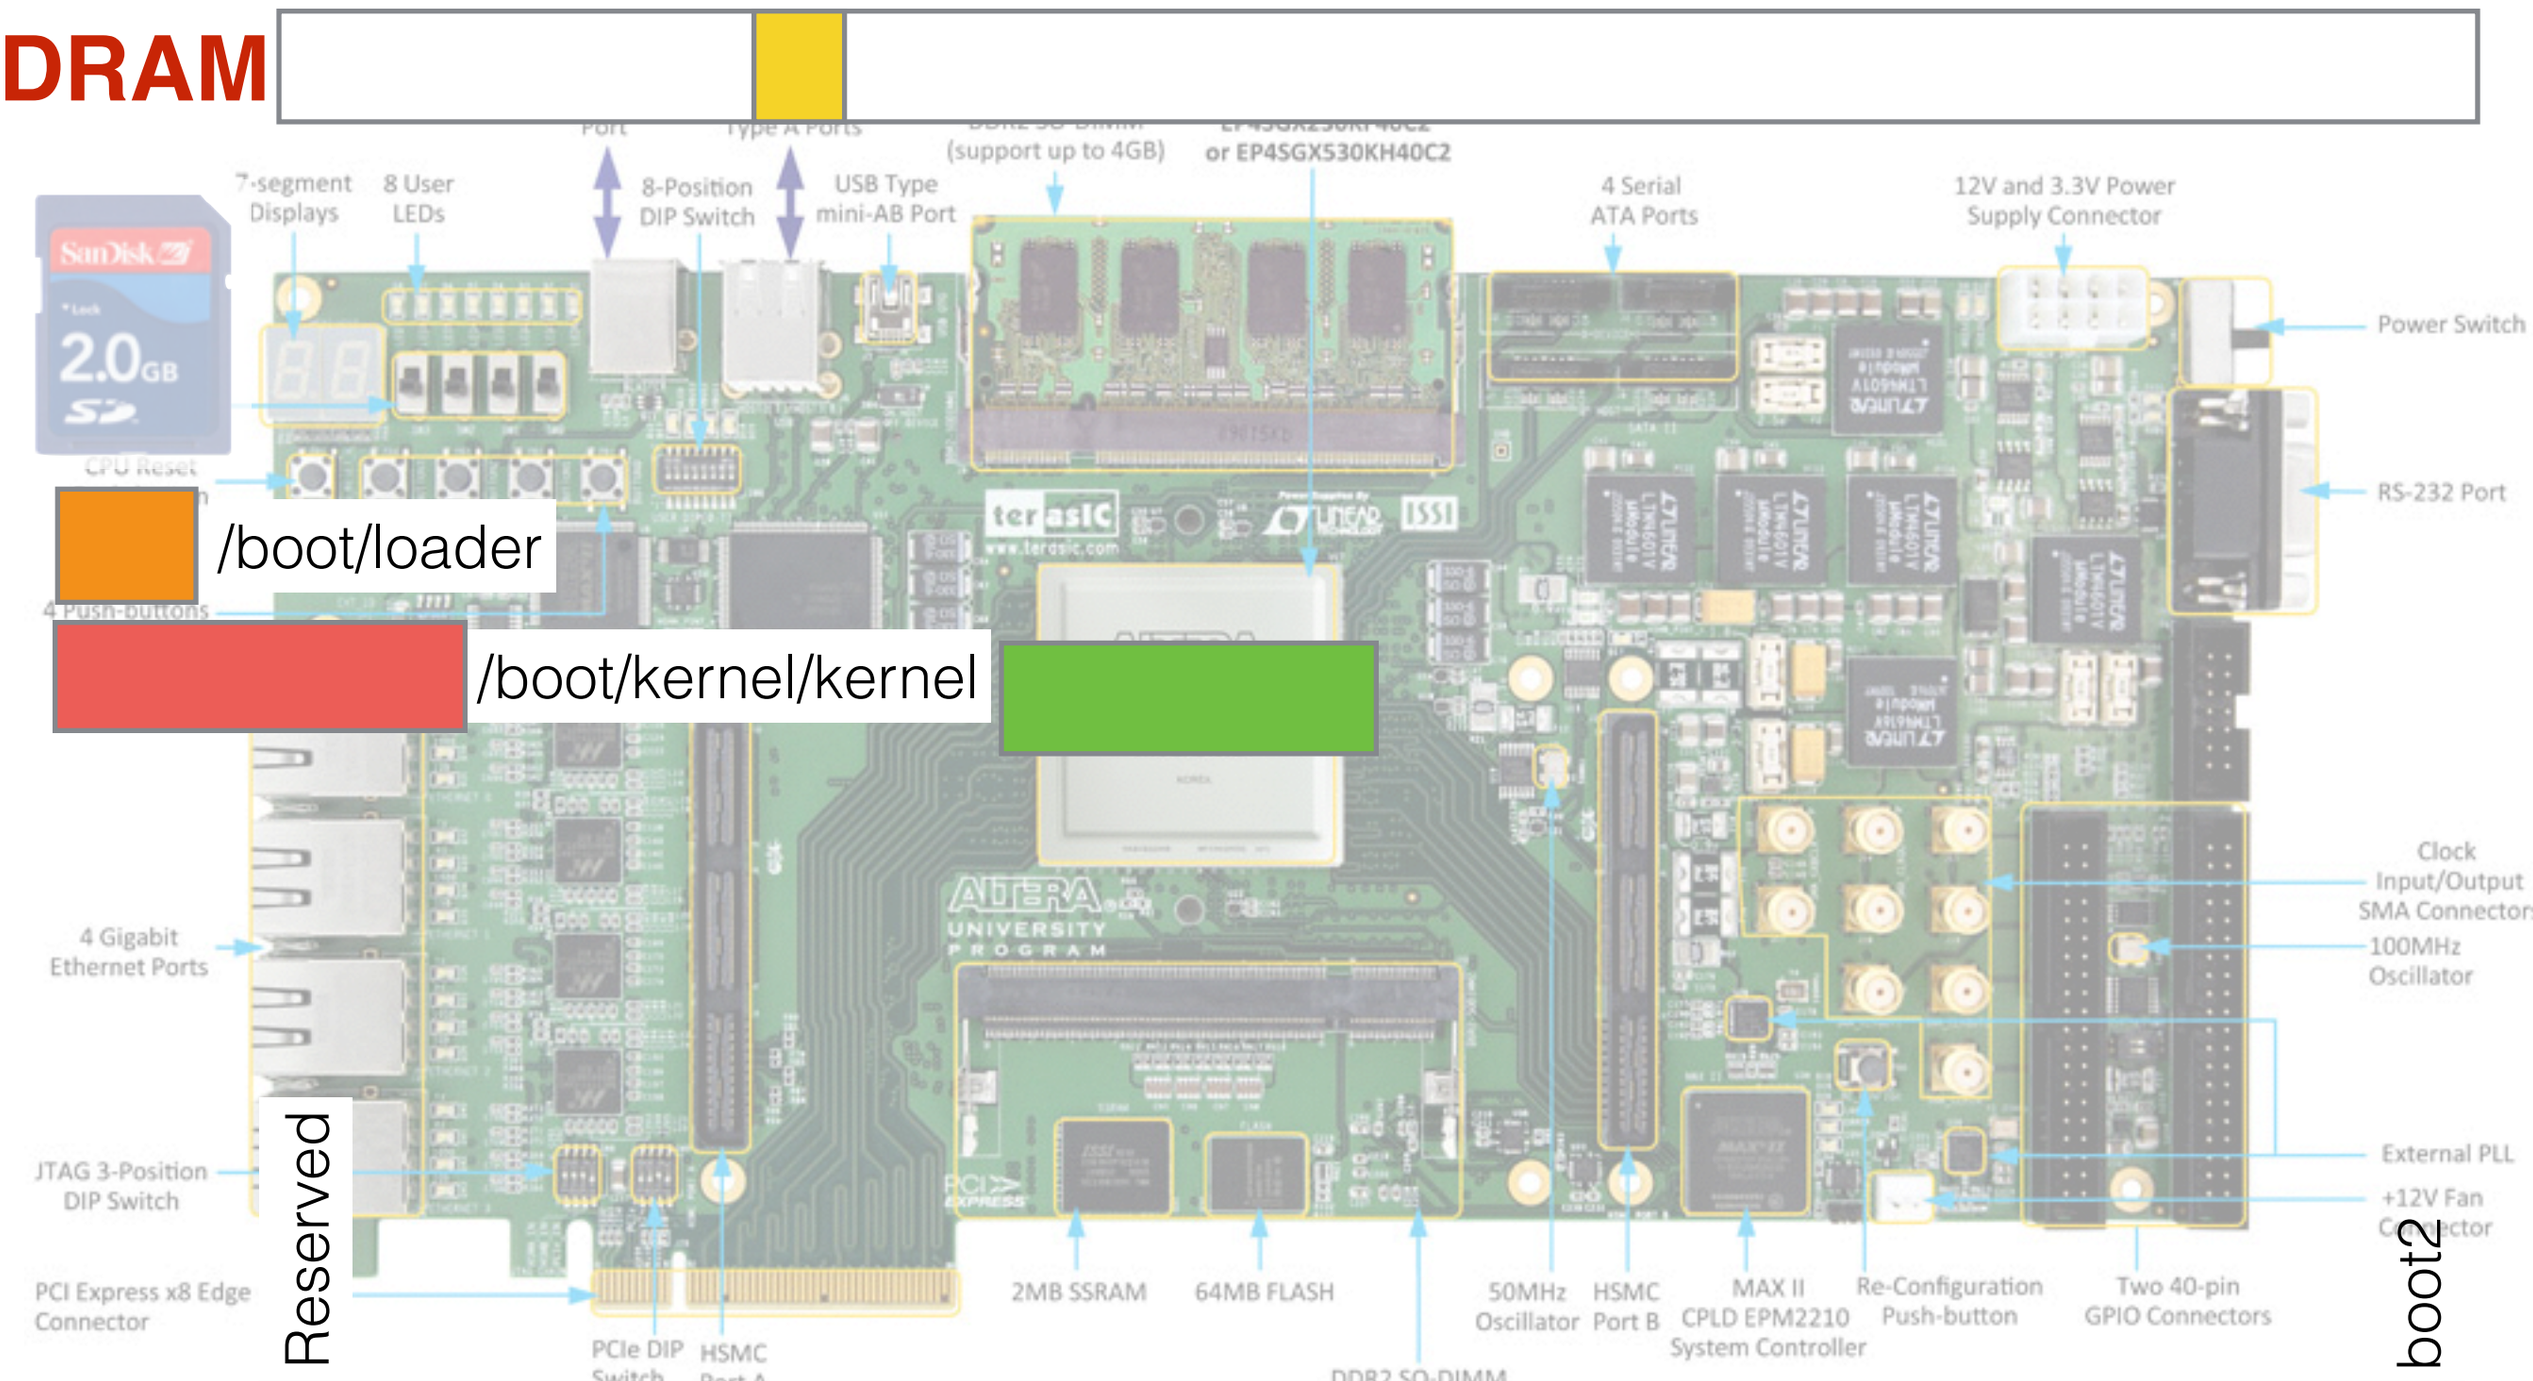

/boot/loader

/boot/kernel/kernel

Flash

	FPGA 1	FPGA2	Operating system (UFS)	
128k	12MB	12MB	39.75MB	128k

Booting the DE4

DRAM

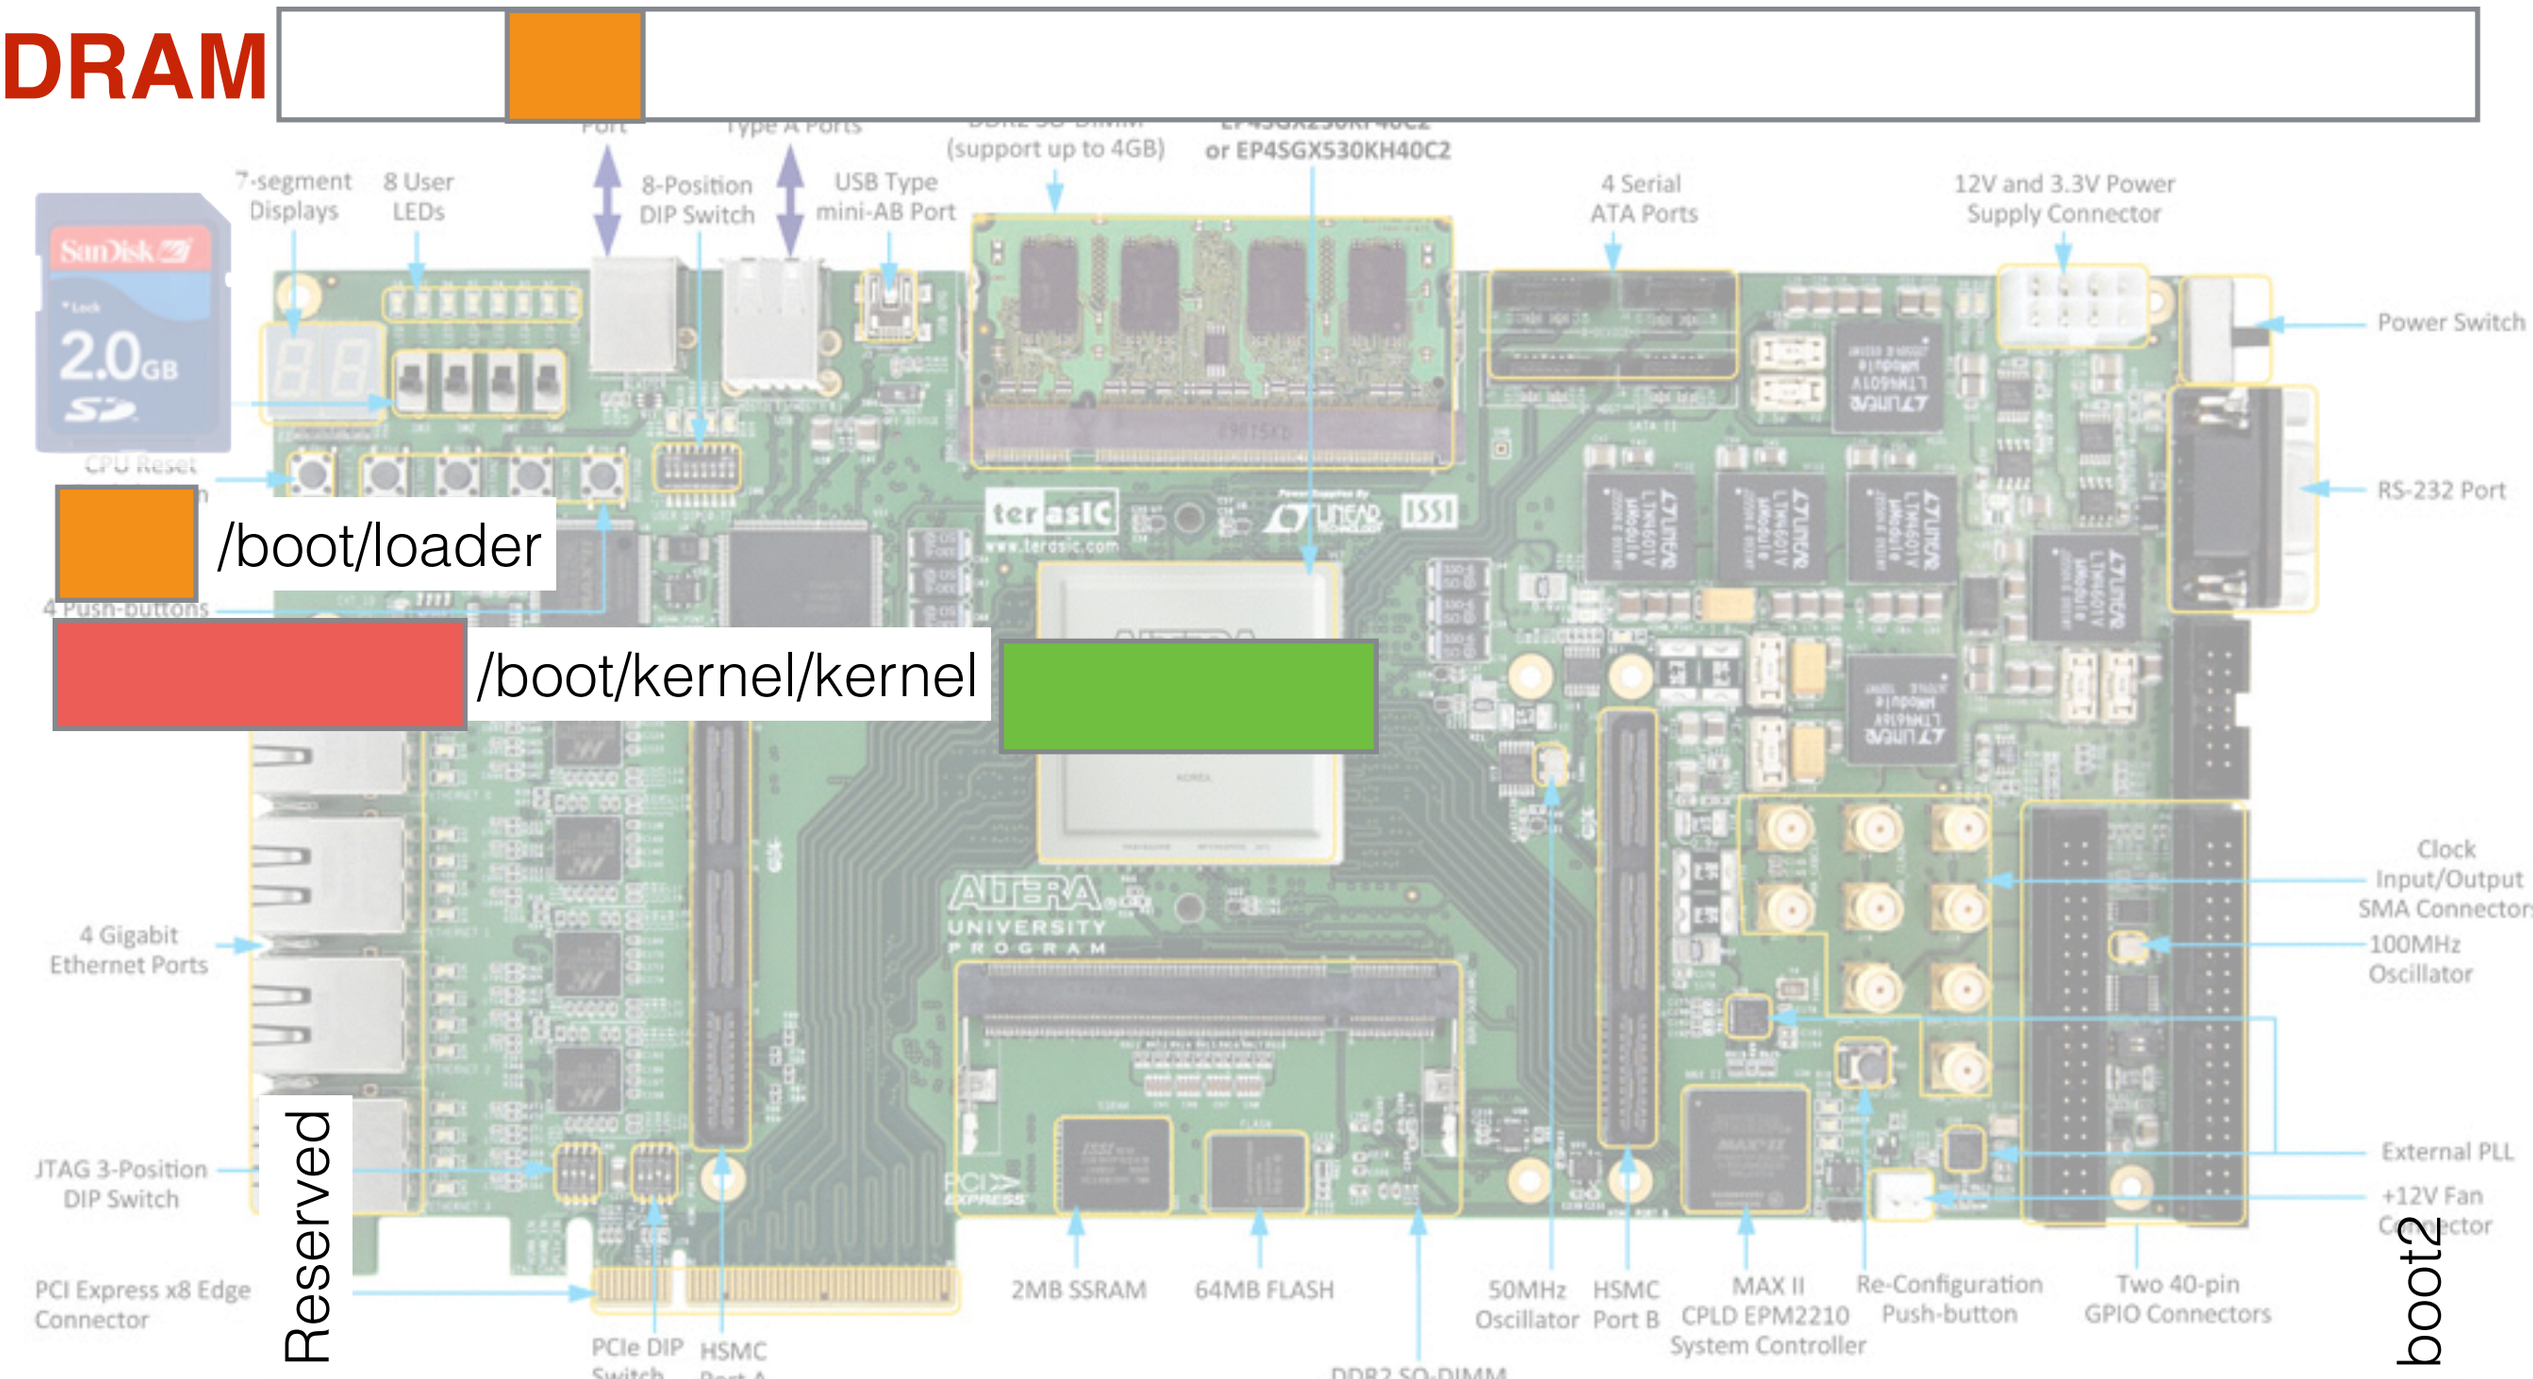

Flash

	FPGA 1	FPGA2	Operating system (UFS)	
128k	12MB	12MB	39.75MB	128k

Booting the DE4

DRAM

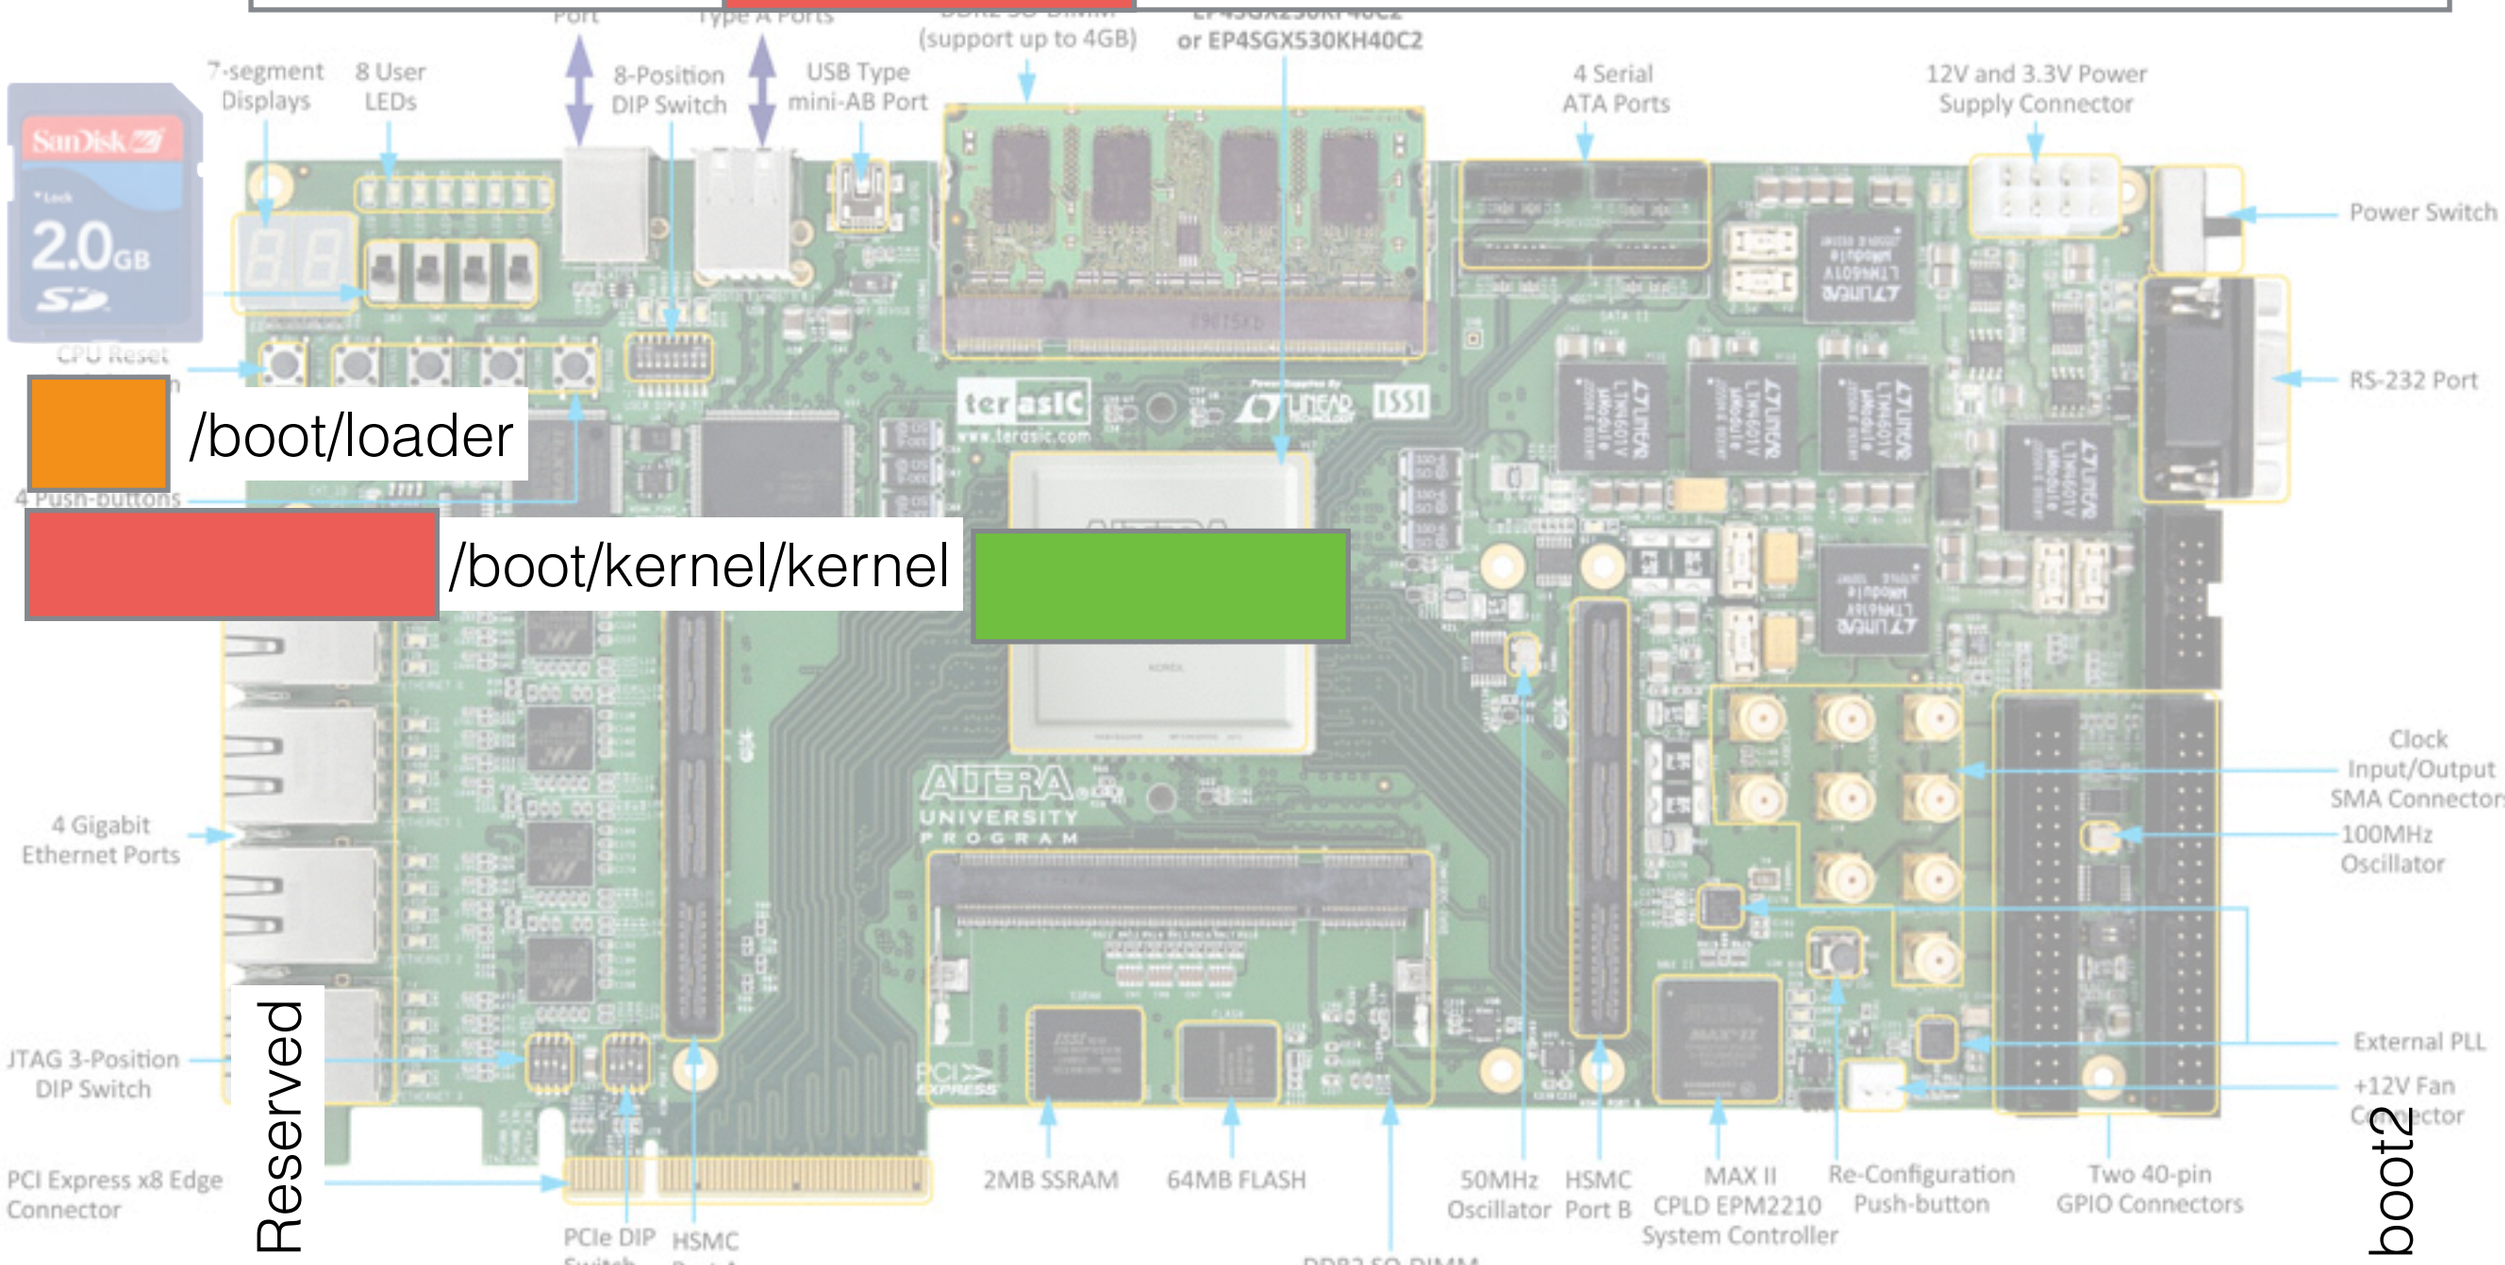

/boot/loader

/boot/kernel/kernel

Flash

	FPGA 1	FPGA2	Operating system (UFS)	
128k	12MB	12MB	39.75MB	128k

Early kernel boot

- Enter kernel in `_locore()` at `_start`
- Calls BERI specific `platform_init()`
- Finally calls `mi_startup()`

SYSINITs

- Initializer functions
- Declared with `SYSINIT()` macro
- Sorted by *subsystem* and *order*
- Run in `mi_startup()`

SI_SUB_COPYRIGHT

SI_SUB_CONFIGURE

- `configure_first()` @ SI_ORDER_FIRST
  - Attaches root bus
- **`configure()`** @ SI_ORDER_THIRD
  - Attaches all devices
  - BERI uses a mix of FDT and hints

Flat Device Tree

```
cpus {
   cpu@0 {
      device-type = "cpu ";
      compatible = "sri-cambridge,beri";
   };
};
soc {
   beripic : beripic@7f804000 {
      compatible = "sri-cambridge,beri-pic”;
      interrupt-controller;
      reg = <0x7f804000 0x400 0x7f806000 0x10
0x7f806080 0x10 0x7f806100 0x10>;
   };
   …
};
```

SI_SUB_EXEC

```
static Elf64_Brandinfo freebsd_brand_info = {
        .brand          = ELFOSABI_FREEBSD,
        .machine        = EM_MIPS,
        .compat_3_brand = "FreeBSD",
        .emul_path      = NULL,
        .interp_path    = "/libexec/ld-elf.so.1",
        .sysvec         = &elf64_freebsd_sysvec,
        .interp_newpath = NULL,
        .flags          = 0
};

SYSINIT(elf64, SI_SUB_EXEC, SI_ORDER_ANY,
    (sysinit_cfunc_t) elf64_insert_brand_entry,
    &freebsd_brand_info);
```

Starting /sbin/init

- `create_init()` at `SI_SUB_CREATE_INIT`
  - Creates process
- `kick_init()` at `SI_SUB_KTHREAD_INIT`
  - Makes process runnable
- `scheduler()` at `SI_SUB_RUN_SCHEDULER`
  - Schedules processes

SMP

SI_SUB_TUNABLES

Sets `mp_cpus` and `mp_maxid` from `platform_cpu_mask()`

SI_SUB_CPU

BSP **AP*n***

```
start_ap(n) {
   …
   cpus = mp_naps;
   platform_start_ap(n);
```

SI_SUB_CPU

BSP **AP*n***

```
start_ap(n) {
   …
   cpus = mp_naps;
   platform_start_ap(n);
                                  mpentry() →
   while (mp_naps <= cpus)
     DELAY(1000);
```

SI_SUB_CPU

BSP

```
start_ap(n) {
   …
   cpus = mp_naps;
   platform_start_ap(n);

while (mp_naps <= cpus)
   DELAY(1000);
```

AP*n*

```
mpentry() →
smp_init_secondary() {
   …
```

SI_SUB_CPU

BSP

```
start_ap(n) {
   …
   cpus = mp_naps;
   platform_start_ap(n);

while (mp_naps <= cpus)
   DELAY(1000);
```

AP*n*

```
mpentry() →
smp_init_secondary() {
   …
   mp_naps++;
   …
   while (!aps_ready)
      ;
```

SI_SUB_CPU

BSP

```
start_ap(n) {
   …
   cpus = mp_naps;
   platform_start_ap(n);

   while (mp_naps <= cpus)
      DELAY(1000);
}
```

AP*n*

```
mpentry() →
smp_init_secondary() {
   …
   mp_naps++;
   …
   while (!aps_ready)
      ;
```

Spin table

```
struct {
        uint64_t entry_addr;
        uint64_t a0;
        uint32_t rsvd1;
        uint32_t pir;
};
```

SI_SUB_SMP

BSP

AP*n*

```
smp_init_secondary() {
  …
  while (!aps_ready)
;
```

SI_SUB_SMP

BSP

```
release_aps() {
/* IPI setup */
…
```

AP*n*

```
smp_init_secondary() {
  …
  while (!aps_ready)
        ;
```

SI_SUB_SMP

BSP

```
release_aps() {
 /* IPI setup */
 …
 aps_ready = 1;

while (!smp_started)
   ;
```

AP*n*

```
smp_init_secondary() {
  …
  while (!aps_ready)
                ;
```

SI_SUB_SMP

BSP

```
release_aps() {
  /* IPI setup */
  …
  aps_ready = 1;

while (!smp_started)
   ;
```

AP*n*

```
smp_init_secondary() {
  …
  while (!aps_ready)
    ;
…
if (/* last AP */)
  smp_started = 1;
```

SI_SUB_SMP

BSP

```
release_aps() {
  /* IPI setup */
  …
  aps_ready = 1;

while (!smp_started)
    ;
```

AP*n*

```
smp_init_secondary() {
  …
  while (!aps_ready)
    ;
…
if (/* last AP */)
  smp_started = 1;

while (!smp_started)
  ;
```

SI_SUB_SMP

BSP

```
release_aps() {
  /* IPI setup */
  …
  aps_ready = 1;

while (!smp_started)
    ;
```

AP*n*

```
smp_init_secondary() {
  …
  while (!aps_ready)
    ;
  …
  if (/* last AP */)
    smp_started = 1;

  while (!smp_started)
    ;
  …
  /* enter scheduler */
}
```

```
last pid:   550;  load averages:  1.93,  0.88,  0.38    up 0+00:03:42  20:16:35
7 processes:   4 running, 3 sleeping
CPU: 69.9% user,  0.0% nice,  2.7% system, 27.4% interrupt,  0.0% idle
Mem: 28M Active, 5624K Inact, 14M Wired, 5440K Buf, 898M Free
Swap:

  PID USERNAME    THR PRI NICE   SIZE    RES STATE    TIME   WCPU COMMAND
  543 root          1 102    0 37996K 28888K RUN      0:46 100.00% pictview
  550 root          1  20    0 11936K  2600K RUN      0:01  0.98% top
  545 root          1  52    0 10476K  1928K RUN      0:01  0.00% getty
  546 root          1  52    0 10476K  1928K RUN      0:01  0.00% getty
  524 root          1  20    0 10488K  1484K select   0:00  0.00% syslogd
  547 root          1  20    0 11104K  2588K wait     0:00  0.00% sh
  544 root          1  52    0 10476K  1928K ttyin    0:00  0.00% getty
```

http://beri-cpu.org

bringing MIPS

Porting FreeBSD

FreeBSD Journal
http://freebsdjournal.org

to a new CPU, even within a previously supported family, is a significant undertaking.