Fault-Tolerant Systems	Reliable Systems
Failure Semantics	reliability
Eault-Tolerant software	• The probability that a system will perform in a correct manner
Fault-Tolerant bardware	availability
Fault-Tolerant data storage	• The probability that a system will deliver a correct answer at one moment
Pessimistic solution — Stable storage	• calculated from the system <i>reliability</i> and expected repair time after a fault.
<ul> <li>Ontimistic solution — Replicated data and network partitioning</li> </ul>	□ fault-avoidance
Physical Clocks	• To try and make a perfect system
O MARS algorithm	□ fault-tolerance
The most unreliable environment	• To try and cope with errors using redundant parts or measures.
Departing Concerls Dashlare	
Byzantine Generals Problem	
(65) - DISTRIBUTED SYSTEMS Fault-Tolerant Systems - Sven Ame Andreasson - Computer Science and Engineering	2 (65) - DISTRIBUTED SYSTEMS Fault-Tolerant Systems - Sven Ame Andreasson - Computer Science and Engineering
CHALMERS	CHALMER
Fault-Tolerance	Fault-Tolerance
	methods
A system might (for certain) contain faults.	□ redundancy
• Programming error	• hardware
• Hardware error	○ software
• Data error	○ data
□ When the execution of the system program hits a fault we will have an error.	• operation
• Then the wanted output of the program can not be achieved in the normal manner.	Different ambition
• If we have implemented measures to cope with this error we will still get a (almost) correct behavior	• error detection and recovery
of the system. This is called fault-tolerance. The error has been masked from the system behavior.	• the errors are visible
When the system "user" sees the error or if the error leads to an inconsistent state or behavior then we have	• error masking
a failure.	• the errors are hidden
	• gives transparency to the errors
65) - DISTRIBUTED SYSTEMS Fault-Tolerant Systems - Syen Ame Andreasson - Computer Science and Engineering	4 (65) - DISTRIBUTED SYSTEMS Fault-Tolerant Systems - Sven Ame Andreasson - Computer Science and Engineering



## **Failure Semantics (2)**

### **Failure Semantics (3)**



- Synchronous Value Fault All responses have the same Consistent Value Fault.
- Consistent Commission Fault All responses have Consistent Timing Fault and Consistent Value Fault
- Commission Fault Responses have arbitrary faults, anything might be wrong in any manner.
  - arbitrary fault
  - Byzantine fault





() CHALMERS

12 (65) - DISTRIBUTED SYSTEMS Fault-Tolerant Systems - Sven Arne Andreasson - Computer Science and Engineering





Recovery Blocks	Exceptions	
Make two different implementations of critical code.		
Use just one of them.	□ Test output values and use an exception if they are not within bounds.	
□ Include a test for each output that it is within reasonable bounds.	Catch exceptions and try to get a result in another way.	
• If not, use the other implementation		
7 (65) - DISTRIBUTED SYSTEMS Fault-Tolerant Systems - Sven Ame Andreasson - Computer Science and Engineering	18 (65) - DISTRIBUTED SYSTEMS Fault-Tolerant Systems - Sven Arne Andreasson - Computer Science and Engineering	
Fault-Tolerant hardware	 Fault-Tolerant data storage	
Use redundant hardware with the same software:		
• comparing pairs	□ Redundancy:	
use two different hardware units and signal an error when values differs	• multiple copies of data objects	
<ul> <li>triple modular redundancy, TMR</li> <li>use the majority value from three different bardware units</li> </ul>	• file level	
	e.g. LOCUS multiple file copies	
Use tests within programs	• block level e.g. stable storage	
• software acceptance test	• disk level	
test that the software has not been corrupted by using a checksum	$\Rightarrow$ inconsistent data	
• multiple units with voting	• we need atomic transactions	
• multiple units with reconfiguration		
• process migration when a computer fails a process might be moved to another computer		
e.g. TANDEM computers		
e.g. TANDEM computers		J

Stable Data Storage		S	Stable Data Stora	ge	
Redundant File System on the Block level - Lampson 1981.					
Instead of calculating the probability of success (which is impossible without reliable input data) we will instead show the different requirements for success.	The Stable Data S	Storage is a layered arcl	hitecture:		
A system consisting of:	Stable Layer	Page copies			Group Error Masking
• client programs					
• data storage units, <i>file servers</i>	"Careful" Layer	"Careful" Storage	Crash - Restart	Ack -Resend	Hierarchical Error Masking
• address $< f,b >$ corresponds to <i>byte</i> b on the file f					
• communication between client programs and <i>file servers</i>	Physical Layer	Hard Disk	Processor	Network	
□ data storage units					-
• read	• At the Ph	ysical Layer we give a	model for the beha	vior of the units.	
• write	• At the "C	Careful" Layer we use re	etry to take care of	temporary errors	3.
• At the Stable Layer we use multiple units to take care of more permanent errors.					anent errors.
21 (65) - DISTRIBUTED SYSTEMS Fault-Tolerant Systems - Sven Ame Andreasson - Computer Science and Engineering	22 (65) - DISTRIBUTED SYSTEMS F	ault-Tolerant Systems - Sven Arne Andreassor	n - Computer Science and Engineerin	g	CHALMERS
The Physical System			Hard Disks		
Model					
Events	A Hard Disk con	sists of			
desired: normal events	• a number of a	addressable pages			
disaster: error events not foreseen, undetected	• each page has	s a block of data and its	status (good,bad)		
	• the operation	s:			
The different algorithms will solve errors but not disasters.	• proced	ure Put (at:Addre	ss;data:Dbloc	ck);	
	• function	on Get(at:Address	s):		
□ The physical units are:	(Statu	es of event:	.a.DD10CK);		
• secondary memory (hard disk)	• result of Dut	or Cot			
• data communication links	result of Put or Get				
• spontaneous events					
23 (65) - DISTRIBUTED SYSTEMS Fault-Tolerant Systems - Sven Ame Andreason - Computer Science and Engineering	24 (65) - DISTRIBUTED SYSTEMS F	ault-Tolerant Systems - Sven Arne Andreassor	1 - Computer Science and Engineerin	g	

# Effects of the different operations

Get(A)					
	value of page A in the file	Get(A) return	s		
desired desired	(good,d) (bad,d)	(good,d) (bad,)			
error disaster disaster disaster	(good,d) (good,d) (bad,d) (good,d)	(bad,) (bad,) (good,) (good,d')	$\leq n_k$ times > $n_k$ times d≠d'		

Put(A,d)		
page A becomes		
desired	(good,d)	
error	unchanged	
error	(bad,d)	

25 (65) - DISTRIBUTED SYSTEMS Fault-Tolerant Systems - Sven Ame Andreasson - Computer Science and Engineering

() CHALMERS

Processors

A processor error will be a crash

 $\equiv$  reset

• The processor state before the crash is lost

#### Spontaneous Events

decay

Divide the units into pair of pages that are not decay related i.e. if there is damage to one page in the pair then the other page in the pair should be expected to be OK.

- If there is a spontaneous change (damage) on a page in the pair then there should not be a spontaneous change on the other page in the pair within a time limit  $T_D$ .
  - Place the two pages in the pair on two different disk units.
     damage to both within the time limit T<sub>D</sub> will become a disaster
     other damage will be expected errors

Spontaneous Events				
error	$(good, d) \rightarrow (bad,)$	only change within $[-T_D, T_D]$		
error	$(bad, d) \rightarrow (good, d)$	spontaneous correction		
disaster	$(good, d) \rightarrow (bad,)$	several changes within $[-T_D, T_D]$		
disaster	$(s, d) \rightarrow (s, d')$	$d\neq d'$ , undiscovered change		

□ We assume that addresses are not corrupted (otherwise disaster) according to the undiscovered change above.

26 (65) - DISTRIBUTED SYSTEMS Fault-Tolerant Systems - Sven Arne Andreasson - Computer Science and Engineering

() CHALMERS

# Communication

### □ messages:

<status:(good,bad);data:Mblock;p:Processor>

- two operations:
  - procedure Send(to:Processor;data:Mblock);
  - function Receive():
     (status:(good,bad);data:Mblock);

		Effects of the diff	ferent operations				Spontaneou	is Events	
		Recei	ive()		For	communicatio	on all errors are spontaneous event	S	
		The message is	Receive() returns						
desired <good,d,p> (good,d)</good,d,p>				Spontaneous Events		is Events			
Ue must all	bw a arbitrary	delay on Receive	(bad, hull)			error error error disaster	a message disappears a message is duplicated a message is damaged undiscovered damage or	$<$ good,d,p> $\rightarrow$ <bad,d',p> <math>&lt;</math>bad,d,p&gt;<math>\rightarrow</math><good,d,p> <math>&lt;</math>good,d,p&gt;<math>\rightarrow</math><good,d',p>d≠d'</good,d',p></good,d,p></bad,d',p>	
							or	<good,d,p>→<good,d,p'>p≠p'</good,d,p'></good,d,p>	
		Sen	nd()						
		call	created message						
	desired	l Send(to,d)	<good,d,to></good,d,to>						
20 ((5)					20.((5)				
29 (05) - DISTRIBUTED 313	EMS Faun-Tolerant Syste	ins - sven Arne Andreasson - Computer Science	e and Engineering		50 (05) - DISTI	XIBUTED STSTEMS Fa	un-roterant Systems - Sven Arne Antreasson - Computer Science a	ia Engineering	
		A Stable	e System		fı	unction Ca	arefulGet(at:address):(	Status,Dblock);	
First we def	ine virtual unit	s with less shortcomings	s than the physical units:		be	egin for i := 1	1 to $n_k$ do		
• careful r	nemorv				1	begin	$d_{a+a}$ . Cot $(a+)$		
						if status,	s = good then return (s	status,data)	
• Operations:						end; return(s	tatus,data);		
Careful Repeats At most	<i>Get</i> <i>Get</i> until one j n <sub>k</sub> times.	page which is good is ret	turned.		e	nd;			
• Carefull Repeats	Put Put and then (	<i>Get</i> to check the result u	ntil <i>Get</i> returns <i>good</i> and with sam	value as <i>Put</i> was		cocedure (	CarefulPut(at:Address;d	ata:Dblock);	
suppose	d to write.		0		be	egin	had		
At most	n <sub>k</sub> times.					i := 0;	Dau,		
						<b>while</b> ( st <b>begin</b>	tatus = <b>bad</b> ) <b>and</b> ( i <	$n_{ m k}$ ) do	
□ careful memory will mask occasional errors on a disk unit and transmission errors.				Put(at,d (status,	ata); d) := Get(at); data <b>then</b> status := <b>ba</b> c	a :			
						i := i+1	;	-,	
						end;			



() CHALMERS

### **Stable Storage Stable Storage (cont.)** Consists of stable pages. □ Stable Page: function StableGet(at:Address) : Dblock begin • one data block (ok,data) := CarefulGet(at copy 1); (without status) if not ok then (ok, data) := CarefulGet(at copy 2); • one ordered pair of *Careful*-pages that are **not** decay related. return(data) (every word in the definition is significant!) end; • Operations: procedure StablePut(at:Address;data:Dblock); ○ StablePut begin CarefulPut(at copy 1, data); CarefulPut(at copy 2, data); ○ StableGet end; ○ Cleanup □ The value of a Stable Page will be: • the value of the first page in the pair if its status is *good* • O otherwise the value of the second page in the pair without taking its status in account. 33 (65) - DISTRIBUTED SYSTEMS Fault-Tolerant Systems - Sven Arne Andreasson - Computer Science and Engineering 34 (65) - DISTRIBUTED SYSTEMS Fault-Tolerant Systems - Sven Arne Andreasson - Computer Science and Engineering () CHALMERS () CHALMERS Cleanup **Reliable Communication** procedure Cleanup; begin □ Is achieved according to normal computer communication methods.: for all pages do begin • computer communication protocol (ok1,d1) := CarefulGet(at copy 1);

Cleanup is performed:

end; end;

- at system start
- after each crash
- at least every  $T_D$  seconds

(ok2,d2) := CarefulGet(at copy 2); if (ok1,d1) = (ok2,d2) then skip else if ok1 and not ok2 then CarefulPut(at copy 2,d1) else if not ok1 and ok2 then CarefulPut(at copy 1,d2) else if ok1 and ok2 then CarefulPut(at copy 2,d1) else skip /katastrof/;



### **Explanation of the Partitioning Graph**

The Partitioning Graph for a file F describes which copies that are kept for the file and how the network might be divided into different partitions and if there is performed an update of the file within the partition. A + indicates that there has been at least one update within the partition.

- In our example we start with all copies in the same partition (1). This means that they can communicate with each other and that subsequently the files copies are equal.
- Then the network is divided into two partitions, 2 and 3. Copies A, B and C are in partition 2 and they perform an update which than is based on the file value in partition 1. Copies A, B and C will have the new value.

Copy D is in partition 3 and keeps the value from 1.

- Then the network is divided again. A and B are in partition 4 while C is in partition 5. In partition 4 there is an update based on the value in partition 2. A and B gets the new value while C keeps the value from partition 2.
- Then part of the network is repaired so C and D gets connected in partition 6. Then D has the old value from partition 1 while C has the value from partition 2. Since C's value is a newer value based on the value D has there is no real conflict, It is just to replace D's old value with C's new value and then they have the same value.
- Then the + for partition 6 indicates that an update. It is based on the value in partition 2.

**Data Consistency (1)** 

• Then the network is completely repaired so all copy servers can communicate. Then they will have two conflicting values, both based on value in partition 2.

```
41 (65) - DISTRIBUTED SYSTEMS Fault-Tolerant Systems - Sven Arne Andreasson - Computer Science and Engineering
```

Pessimistic methods

o primary copy

*token*voting

disallow concurrent updating/reading

might give performance/availability problems.

() CHALMERS

42 (65) - DISTRIBUTED SYSTEMS Fault-Tolerant Systems - Sven Arne Andreasson - Computer Science and Engineering

CHALMERS





The main problem is to even avoid such conflicts to happen, pessimistic methods, or detect when they have happened.

Data Consistency (2)	Version Vectors
	□ Version Vector for a file:
Optimistic methods	$\bigcirc$ a sequence of <i>n</i> pairs
• allow updates without locking	<i>n</i> number of copies of the file
• solve conflicts when partitions are merged <i>reconciliation</i>	pair <i>i</i> : $(S_i : v_i)$ in the sequence indicates the number of updates that have been made at the node $S_i$ keeping a copy of the file
• Roll Back	• Each copy of the file keeps a corresponding Version Vector.
cheap updates	$\Rightarrow$ Version Vectors are the same as Vector Clocks
• expensive <i>reconciliation</i>	□ When updating a file;
<ul> <li>detection of conflicts:</li> <li>Detects all possible conflicts without false alarms.</li> <li>Version Vectors = Vector Clocks</li> </ul>	<ul> <li>Each update has a responsible node which has one of the file copies. The update is performed based on this file copy and the new version gets a Version Vector that is the old copy's Version Vector with the nodes own value interconnected by 1. Then all reachable copies of the file get the new version of the file together with the new Version Vector.</li> <li>Multiple updates within one partition is not allowed, there can only be conflicts between updates in different partitions.</li> </ul>
45 (65) - DISTRIBUTED SYSTEMS Fault-Tolerant Systems - Sven Ame Andreason - Computer Science and Engineering	46 (65) - DISTRIBUTED SYSTEMS Fault-Tolerant Systems - Sven Arne Andreason - Computer Science and Engineering
Version Vectors - Conflict Detection	Version Vectors — Example with no conflicts +ABCDE

- $\Box$  When two partitions, with version vectors  $V_i$  and  $V_j$ , are merged the following might hold:
  - $V_i = V_j$  in all positions:

then no change has been made in any of the partitions. No action has to be taken.

•  $V_i < V_j$  according to the vector clock definition:

then there has been at least one change in partition with vector  $V_j$ . The copies in the other partition should then be replaced with the latter version and they will also get its version vector.

 $\bigcirc$  V<sub>i</sub> // V<sub>i</sub> according to the vector clock definition:

then there has been at least one change in both partitions. This means that we have a conflict that must be solved to get a resulting version in the new partition.

The new version will get the version vector where each position has the maximum value from the two old version vectors. Then one is added for the node that solved the conflict (calculated the new version), it actually does an update to the file.

- Generally this conflict solving must be done manually. No algorithm can understand the content of a general file.
- For a special file such as a directory the conflict can be solved automatically if for instance the two updates have been the creation of two files with different names.



ABCDE

7

DE

3

partition Version Vector

<A:1, B:0, C:0, D:0, E:0> A executes the update

 $\langle$ A:1, B:1, C:0, D:0, E:0 $\rangle$  B executes the update

<A:1, B:0, C:0, D:0, E:0> no update, same as 1

<A:1, B:1, C:0, D:0, E:0> no update, same as 2

<A:2, B:1, C:0, D:0, E:0> A executes the update

<A:1, B:1, C:0, D:0, E:0> D,E gets same as 2

<A:2, B:1, C:0, D:0, E:0> C gets same as 6

ABDE

4

+ABDE

6

1

2

3

4

5

6

7

+ABC

2

С

5

CHALMERS

- In the first partition there is one update performed by A. All copies get the new value and the corresponding version vector will become <A:1, B:0, C:0, D:0, E:0> which is stored with all copies.
- 2. In this partition an update is performed by B. Copies A, B and C will be changed and get the version vector <A:1, B:1, C:0, D:0, E:0>.
- 3. In this partition there is no updates and the copies D and E will keep their values.
- 4. When merging copies A and B with D and E their version vectors are compared. It shows that the version vectors for D and E are strictly dominated by the version vectors of A and B. Thus D and E have older copies that not conflict with copies A and B. The older copies are changed to the new version and then also get the corresponding version vector. The version vector will become <A:1, B:1, C:0, D:0, E:0> for all copies.
- 5. In this partition there is no update.
- 6. In this partition an update is performed by A. Copies A, B, D and E will be changed and get the version vector <A:2, B:1, C:0, D:0, E:0>.
- 7. When C is merged with the others it has an older value. It gets the new value and version vector from the others.

49 (65) - DISTRIBUTED SYSTEMS Fault-Tolerant Systems - Sven Arne Andreasson - Computer Science and Engineering

(M) CHALMERS

- In the first partition there is one update performed by A. All copies get the new value and the corresponding version vector will become <A:1, B:0, C:0, D:0, E:0> which is stored with all copies.
- 2. In this partition an update is performed by B. Copies A, B and C will be changed and get the version vector <A:1, B:1, C:0, D:0, E:0>.
- 3. In this partition an update is performed by D. Copies D and E will be changed and get the version vector <A:1, B:0, C:0, D:1, E:0>.
- 4. When merging copies A and B with D and E their version vectors are compared. It shows that there is a conflict. The conflict is solved by B. The version vector for the new copies will have the maximum value in each position from vectors 2 and 3. Then the value for B is incremented since B resolves the conflict. The version vector will become <A:1, B:2, C:0, D:1, E:0> for A,B,D and E copies.
- 5. In this partition there is an update by C based on the copy from partition 2. The version vector will become <A:1, B:1, C:1, D:0, E:0> which is stored with the C copy.
- 6. In this partition an update is performed by A based on the copy from partition 4. The version vector will become <A:2, B:2, C:0, D:1, E:0> for A,B,D and E copies.
- 7. When C is merged with the others there is a conflict again. After solved by E the version vector for the new copy will be <A:2, B:2, C:1, D:1, E:1> which is stored with all copies.



#### 50 (65) - DISTRIBUTED SYSTEMS Fault-Tolerant Systems - Sven Arne Andreasson - Computer Science and Engineering

#### CHALMERS

# **Version Vectors**

- Version Vectors introduces high availability in case of network part since we are allowed to do updates even if we don't have access to all copies (one is enough!)
  - Good if we know that there is a low risk for simultaneously updates for the same file.
  - If there still would be simultaneously updates we are in trouble, but it will be detected.









# **Byzantine Generals — Computer Applications**

Byzantine faults

Requirements:

• At most *m* of *3*\**m*+*1* nodes fail

U Voting

- The process giving its vote is the General.
- For each voting round each process then takes turn as General.
- Used in the *Space Shuttle* control system:
  - 4 computers running the same program compares their results.
  - Will protect the shuttle from errors due to data change due to cosmic rays if there is only damage in one of the computers.
- The *Space Shuttle* also has a fifth computer with the control program written in another language.
  - This can be shifted to manually if there are doubts about the other program.
  - There can then be no shift back to the first program!
  - N-version programming with N = 2.

65 (65) - DISTRIBUTED SYSTEMS Fault-Tolerant Systems - Sven Arne Andreasson - Computer Science and Engineering

CHALMERS