Safety and Security

Ross Anderson

Security vs Safety

- The safety and security communities use similar processes, but different languages
  - Hazard analysis –> safety case
  - Threat model –> security policy
- Safety and security can be engineered top-down in a project, or by a process as a product evolves
- For us, dependability = reliability + security
- Reliability and security are often strongly correlated in practice
- Safety and security are increasingly entangled!

What's different

- Usability
  - Safe usability, safe defaults, skill context
- Liability
  - Product liability is strict, so blame 'error'
- Risk
  - Risk appetites are different with an adversary present
- Some policies change radically
  - Multilevel policies may allow flow downwards only rather than upwards only

Recall Multilevel Integrity

- In lecture 2, we saw the Biba model – data may flow only down from high-integrity to low-integrity
- Example A: medical device calibration / operation
- Example B: electricity / gas / oil distribution
  - Safety: highest integrity level
  - Monitoring and control: next level
  - Enterprise apps (e.g. billing): third level
- This was the inverse of multilevel confidentiality, where data flows up from confidential to secret to top secret, but never down

Recall Shared Control

- In banking, you need two managers to authorize a large money transfer
- Similarly, some safety-critical activations require multiple people or specially coordinated actions
- Extreme example: the management of nuclear weapons (see book chapter for details)
- Mundane example: industrial guillotine that could cut your hand off – activated by a switch each side
- Teamwork – e.g. pilot / copilot landing a plane
- Usability is often the big soft spot with safety!

Safety, security and human behaviour

- Vendors will try to blame accidents on ‘user error’
- Many car crashes in the early years were down to badly designed cars, roads, regulations...
- For the rest, we now also provide seat belts, airbags, passenger cages, crumple zones...
- Compare 1959, 2009 Chevrolets in crash video:
  - A 25mph crash impaled the driver of the 1959 car on the steering wheel. This would have been fatal
  - The 2009 car’s passenger compartment was intact. The driver would have walked away

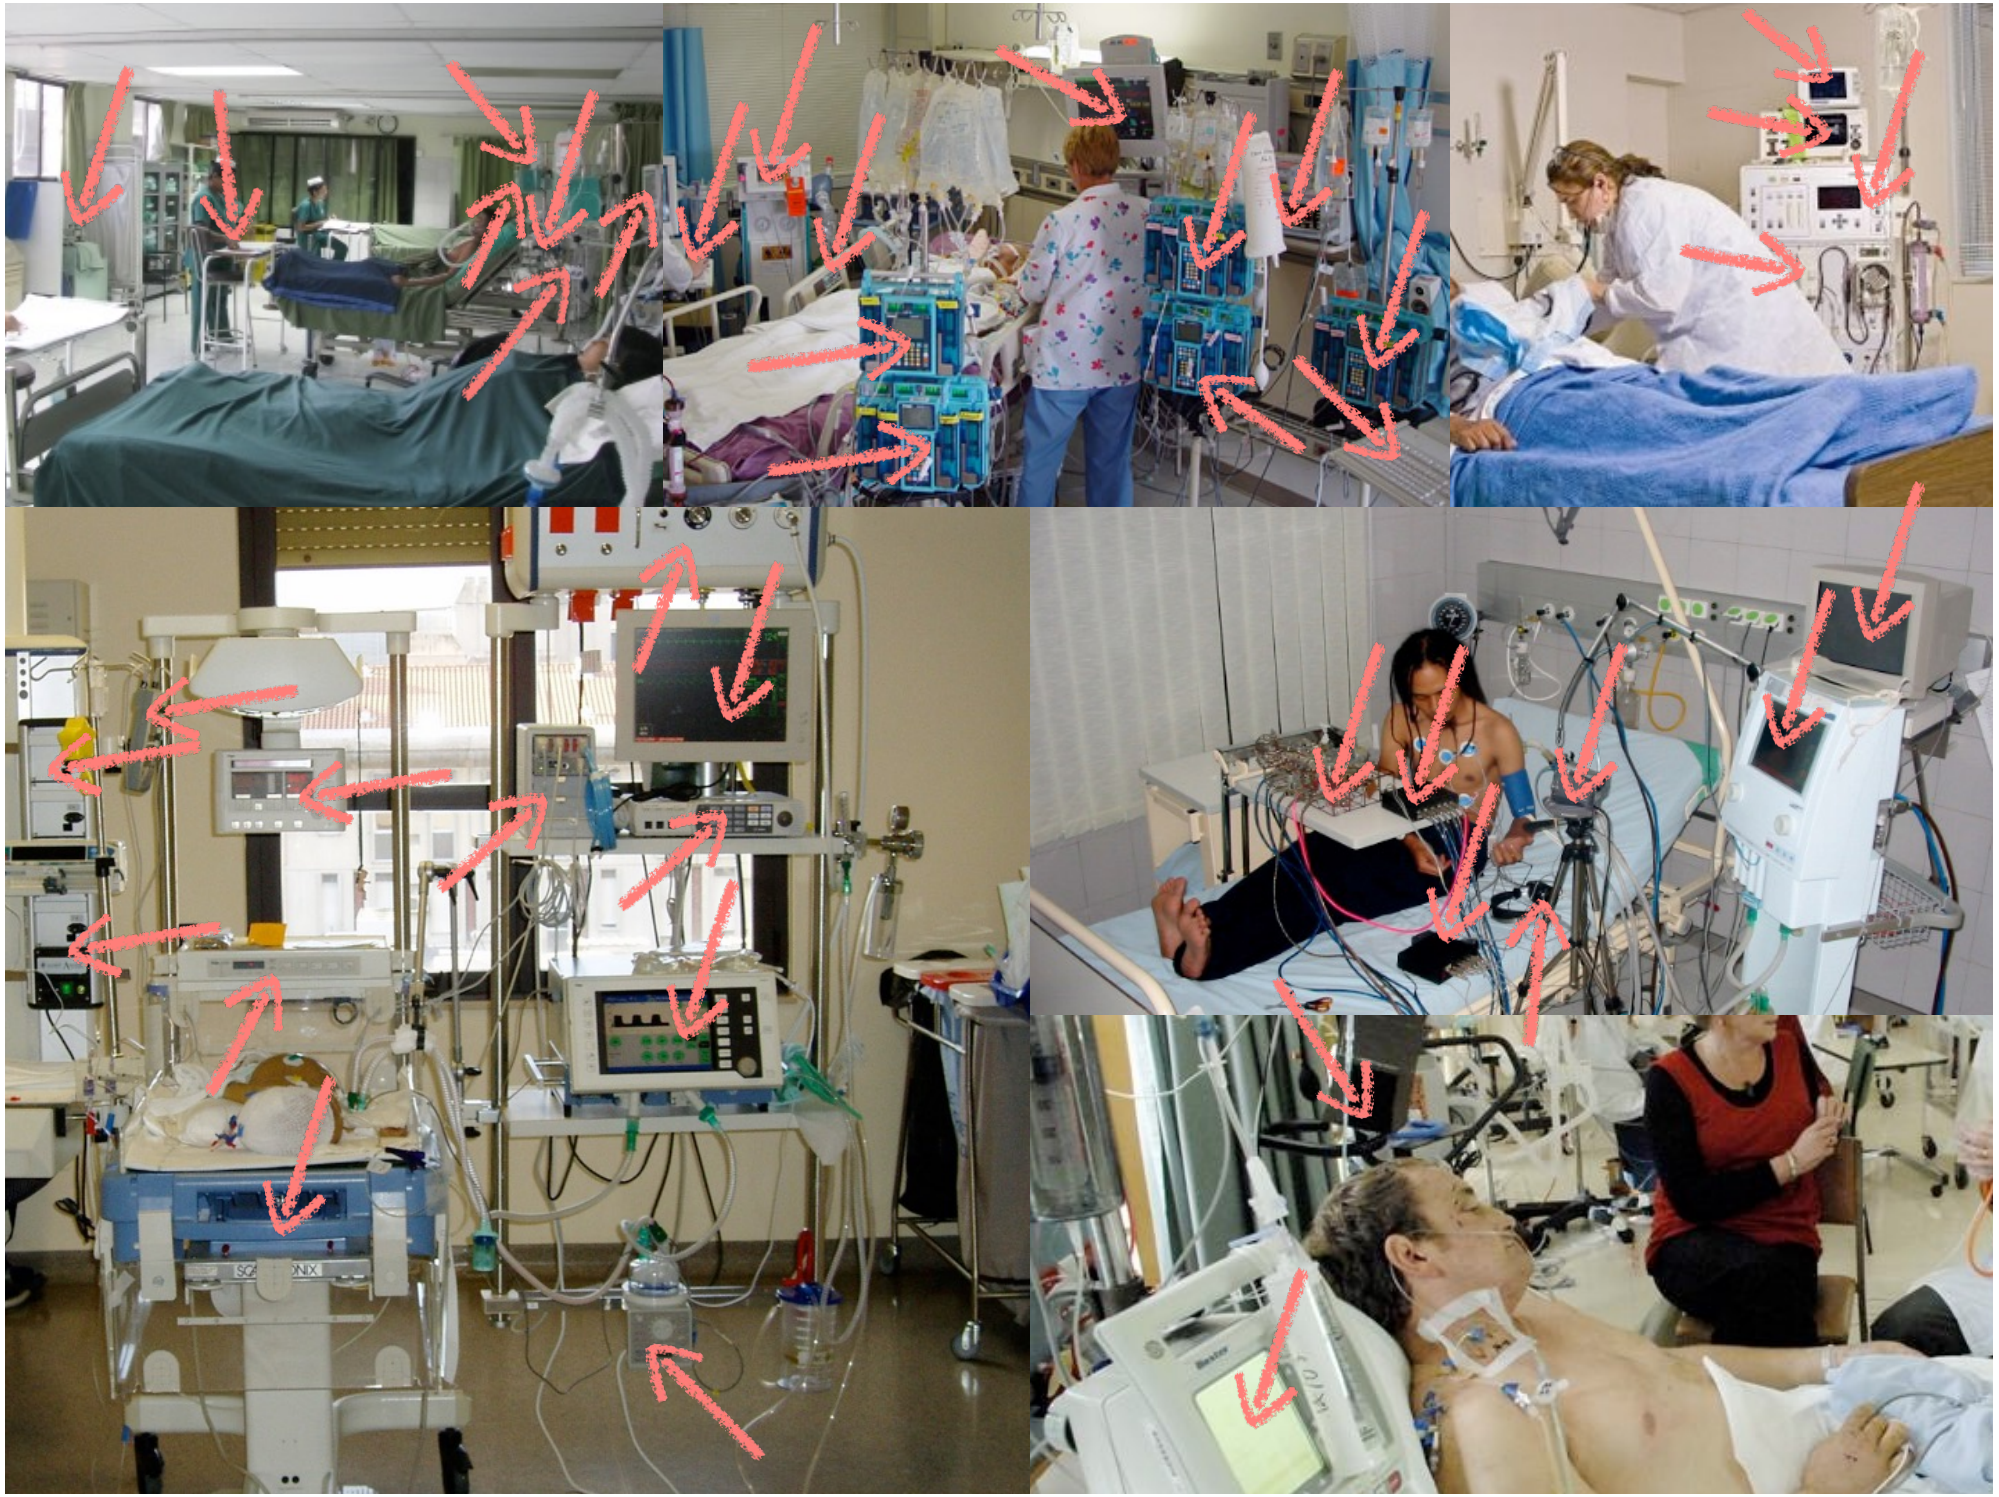


CareFusion
Alaris® PC
PRIMARY INFUSION
RATE
VTBI
175.0
Press START

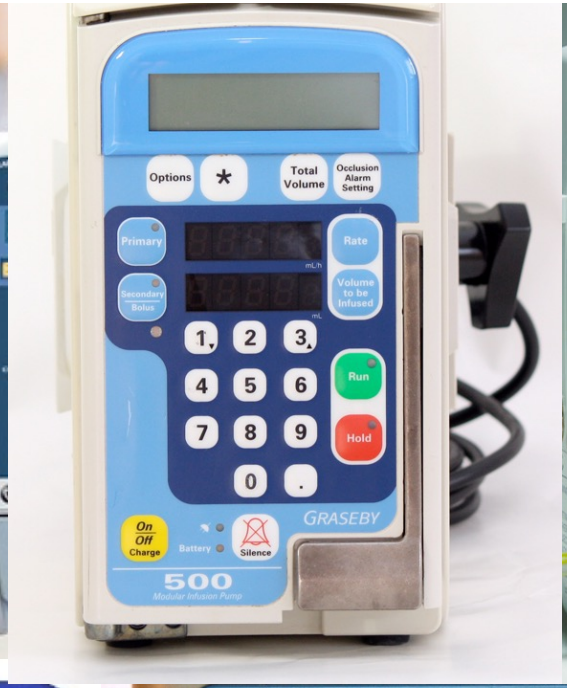
Options
Total Volume
Occlusion Alarm Setting
Primary
Secondary Bolus
Rate
Volume to be Infused
Run
Hold
On Off Charge
Battery
Silence
GRASEBY
500
Modular Infusion Pump

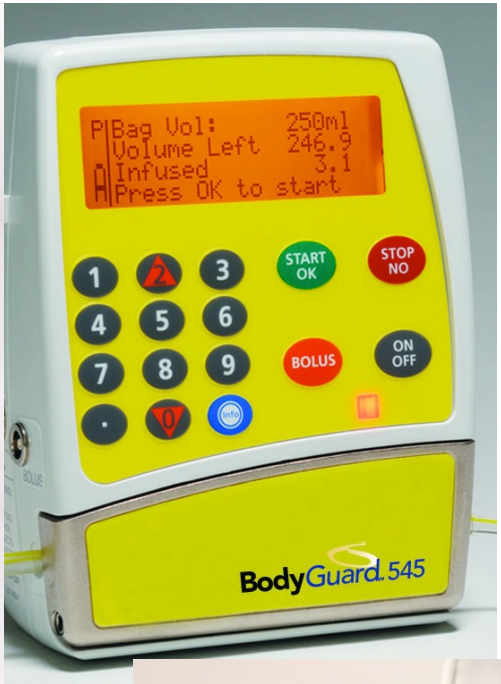
Bag Vol: 250ml
Volume Left 246.9
Infused 3.1
Press OK to start
START OK
STOP NO
BOLUS
ON OFF
BodyGuard 545

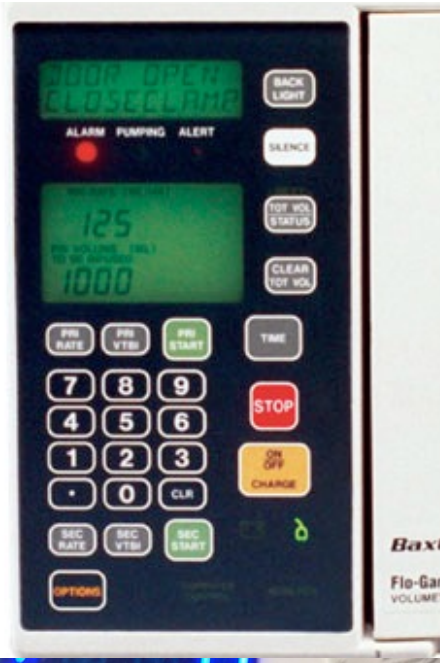
ALARM
PUMPING
ALERT
BACK LIGHT
SILENCE
TOT VOL STATUS
CLEAR TOT VOL
TIME
STOP
ON OFF CHARGE
OPTIONS

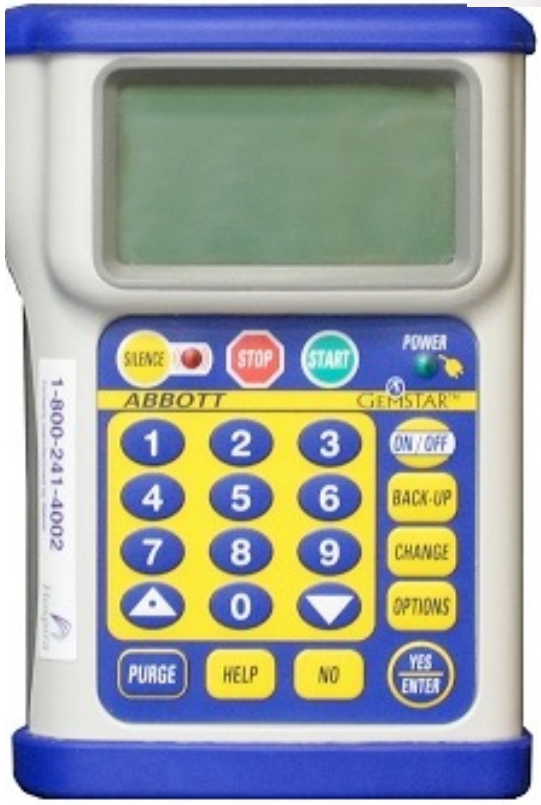
SILENCE
STOP
START
POWER
ABBOTT
GEMSTAR
ON / OFF
BACK-UP
CHANGE
OPTIONS
PURGE
HELP
NO
YES ENTER
1-800-241-4002

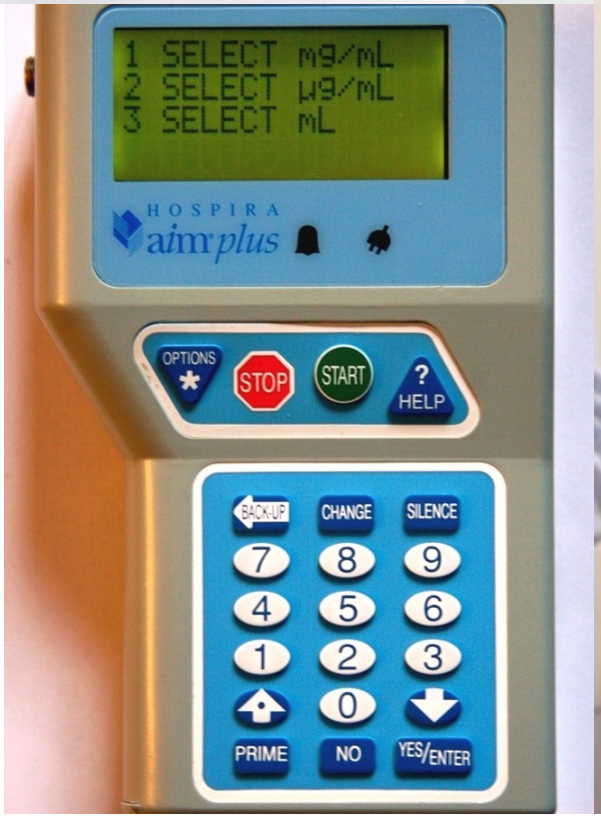
1 SELECT mg/mL
2 SELECT µg/mL
3 SELECT mL
HOSPIRA
aim plus
OPTIONS
STOP
START
HELP
BACK-UP
CHANGE
SILENCE
PRIME
NO
YES/ENTER

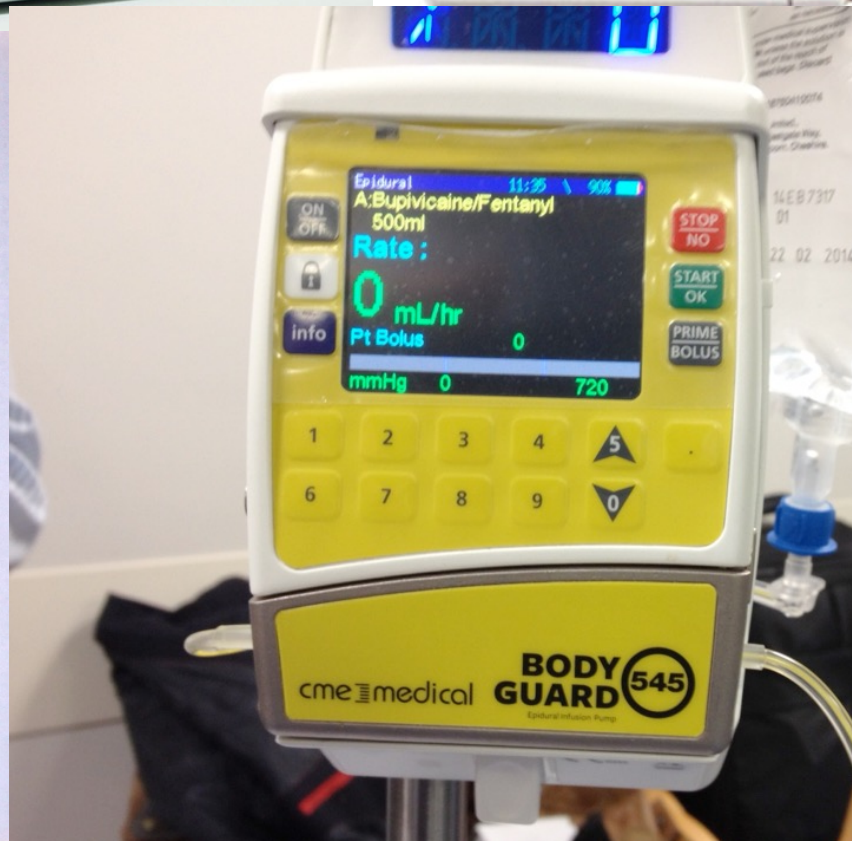
Epidural
A:Bupivicaine/Fentanyl
500ml
Rate :
0 mL/hr
Pt Bolus 0
mmHg 0 720
ON OFF
info
STOP NO
START OK
PRIME BOLUS
cme medical
BODY GUARD 545

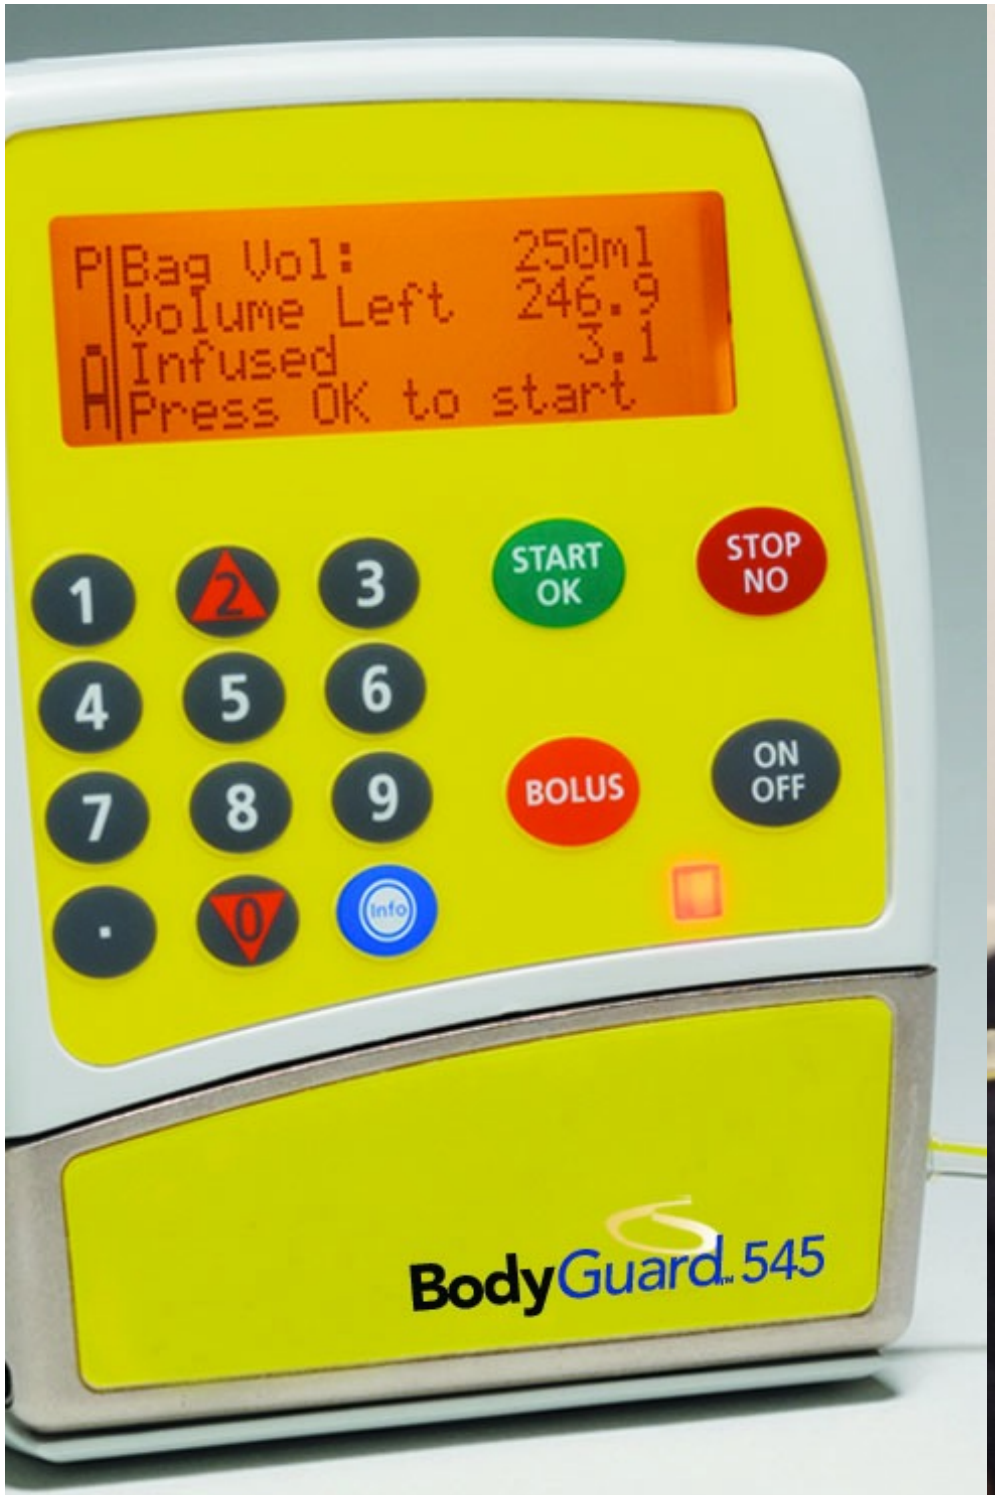
P Bag Vol: 250ml
Volume Left 246.9
A Infused 3.1
Press OK to start
START OK
STOP NO
BOLUS
ON OFF
BodyGuard 545

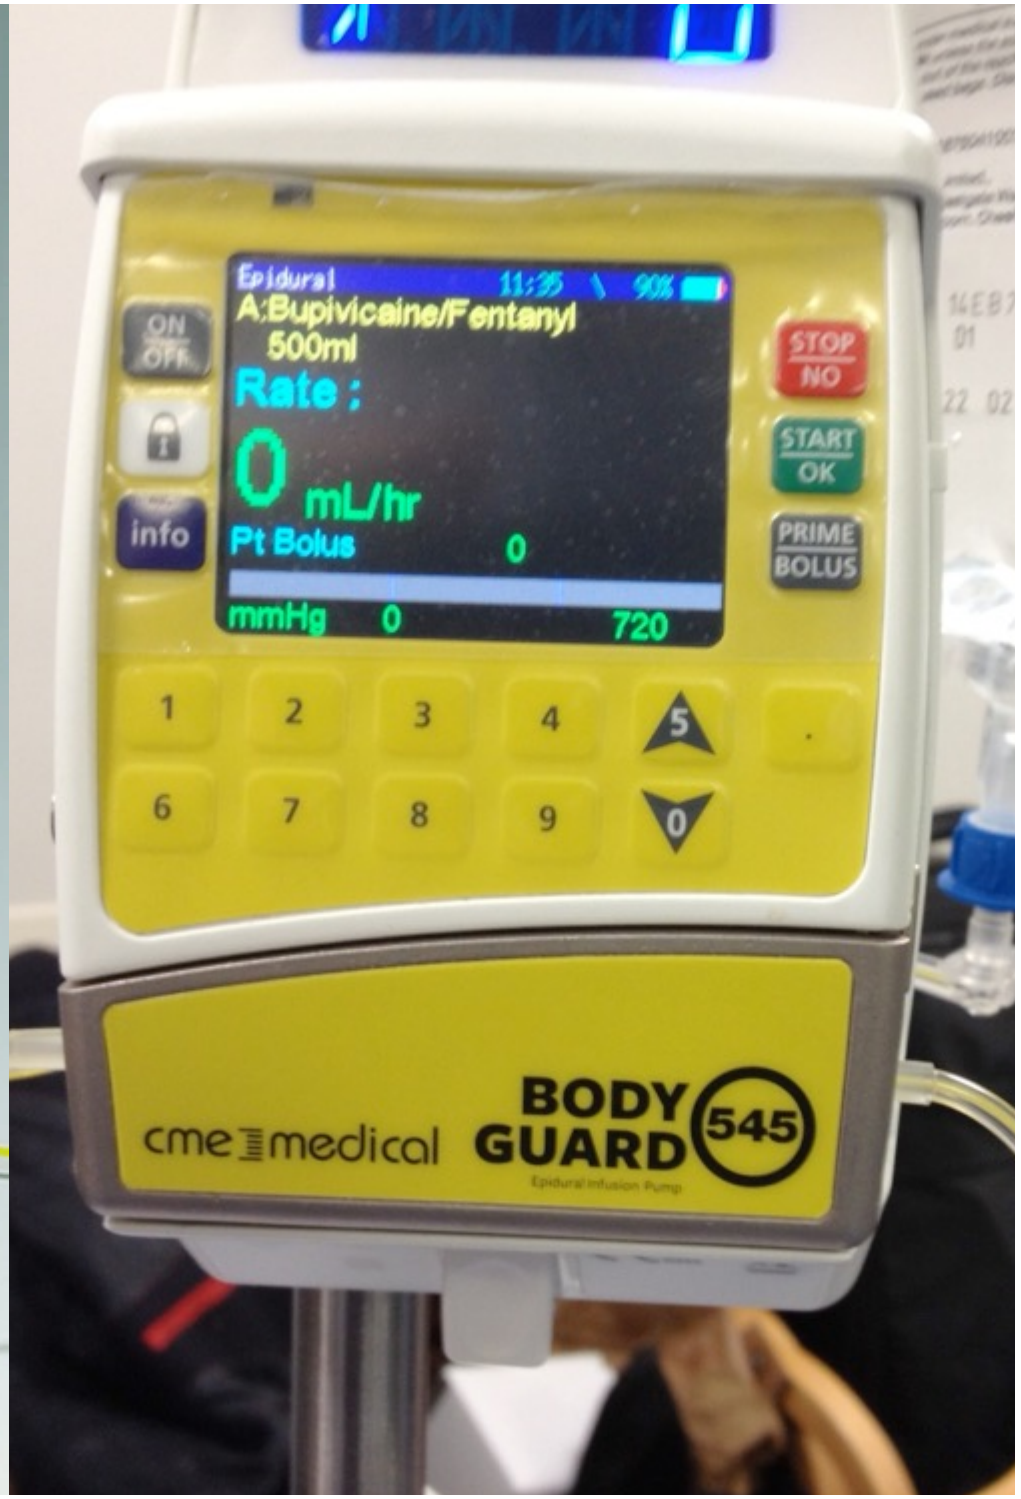
Epidural 11:35 90%
A:Bupivicaine/Fentanyl
500ml
Rate :
0 mL/hr
Pt Bolus 0
mmHg 0 720
ON OFF
info
STOP NO
START OK
PRIME BOLUS
cme medical
BODY GUARD 545

Medical Devices

- Research by Harold Thimbleby: hospital safety usability failures kill about 2000 p.a. in the UK, about the same as road accidents
- Safety usability ignored – incentives wrong…
- But attacks are harder to ignore – Kevin Fu's Wi-Fi tampering demo in 2015 led the FDA to blacklist the Hospira Symbiq infusion pump
- 2017: recall of 450,000 St Jude pacemakers
- 2021: Biden appoints Kevin to the FDA...

Medical Devices (EU)

- The Medical Device Directives have been revised: from 2021 they require post-market surveillance, a per-device risk management plan, ergonomic design …
- Reg 17.2: ‘for devices that incorporate software… the software shall be developed … in accordance with the state of the art taking into account the principles of development life cycle, risk management, including information security, verification and validation’

Progress?

- Security scares have helped to bring attention to safety
- FDA tech staff increased from two to hundreds
- Progress is still a hard slog though!
- Many infusion pumps still have incompatible UIs
- Real problem: entrenched cartel of suppliers
- Same with other medical equipment
- E.g. fatal accidents with medical accelerators since mid-1980s (google ‘Therac 25’) but still happening

Safety case maintenance

- Big issue in many sectors!
- Vendors prefer to front-load certification, whose costs deter new market entrants, and dislike recalls, which are expensive
- Being worked on for medical devices, via post-market surveillance – but we see similar patterns with cars, household goods, even aircraft…
- We'll discuss this systematically in the following lecture on assurance and sustainability

Autonomous vehicles

- Autopilots for aircraft go back to 1914!
- Remotely piloted vehicles (RPVs), Israel/Syria 1982
- 2004: DARPA challenge to build an autonomous vehicle to cover 149 miles of Mojave Desert
- First vehicle to complete was Stanley, from Stanford, in 2005
- Google hired the team in 2009 to work on self-driving cars, announcing the goal in 2010
- Tesla fielded the first product in 2014...

What do we mean by autonomy?

1. Software controls either steering or speed
2. Software controls steering and speed, but driver must take over at zero notice if software fails
3. Software also monitors the environment but assumes human available as fallback
4. Software drives the car without assuming a human can intervene. If it gets confused, it stops at the side of the road
5. Software can do everything a human can

What do we have?

- Automated driver assistance systems (ADAS) operate at levels 1 or 2, where you assume a human driver can take over at zero notice
- Crashes happen when drivers assume ‘autopilot’ means level 4, while they’re getting 2
- But surely level 2 can help? After all, falling asleep at the wheel causes 20% of accidents, 30% of fatal ones (50% on motorway)
- If automatic cruise control plus automatic lane keeping plus automatic emergency braking prevent half of these, surely deaths should fall?

What do we get?

- Tesla claimed fewer fatal accidents when ‘autopilot’ was engaged on the freeway (and initially persuaded NHTSA)
- Independent analysis showed there were more. Why might this be?
  - Risk thermostat – people adapt to a perceived reduction in risk
  - Affordances – ‘nothing to do’, so people relax
  - Industry marketing
  - Autopilot switches off when it gets confused

Why it’s hard

- It takes about eight seconds for a commercial pilot to assume control after an autopilot failure
- With cars, it’s worse; the driver (and in a taxi, the safety driver) may be watching a movie
- One standard for level 4 says that the car must come to a stop in the same lane
- But often cars in roadworks get confused about the lane, so this isn’t even testable
- Even turning across traffic is hard right now...

Even braking is hard!

- Adaptive cruise control: at N mph, stay N yards from the car in front
  - Implementation: radar
- Emergency brake assist: help driver do a hard stop
  - Implementation: detect acc -> brake in 300mS then 2 kg
- Automatic emergency braking: if needed
  - Implementation: various, but very tricky in town traffic
- Do you save more front-end crashes than you cause rear-end crashes? Do you have too many false alarms?
- What if one of these systems can be hacked?

When cars get hacked

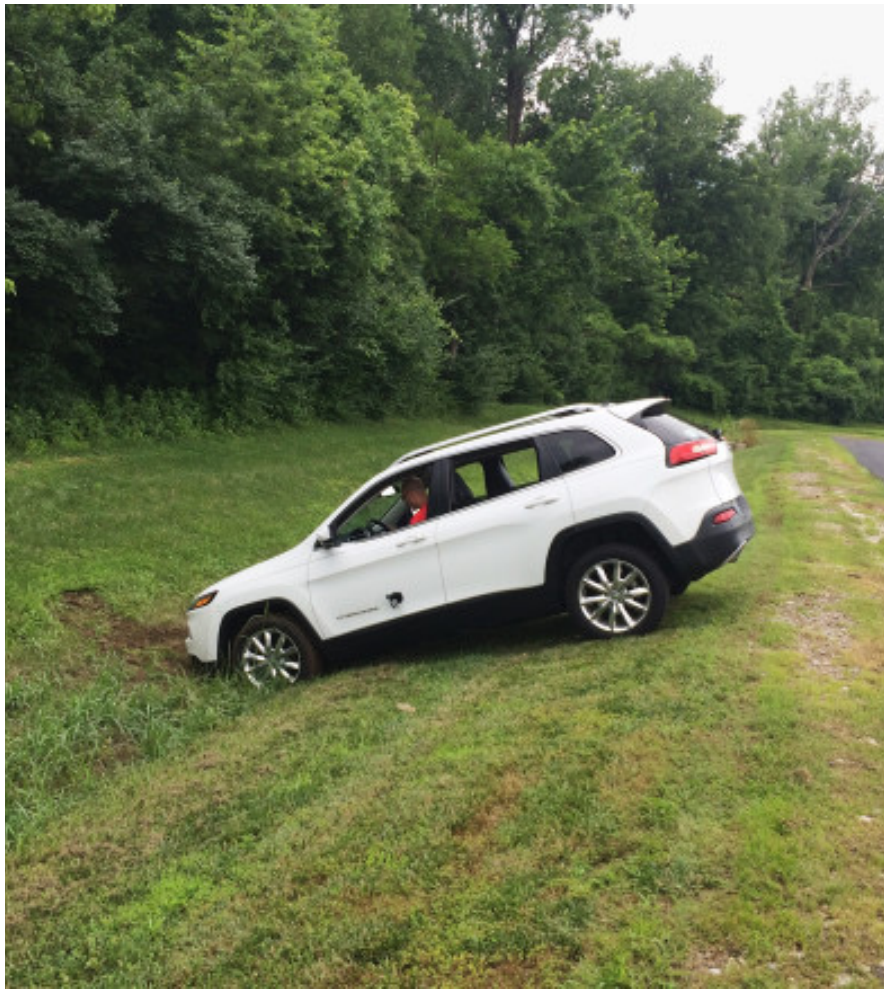

- Old days: ignore hacks or sue critics
- 2015: Charlie Miller and Chris Valasek hacked a Jeep Cherokee over the mobile network
- Suddenly people cared…
- Chrysler recalled 1.4m vehicles for software fix, costing over $1bn

When cars get hacked (2)

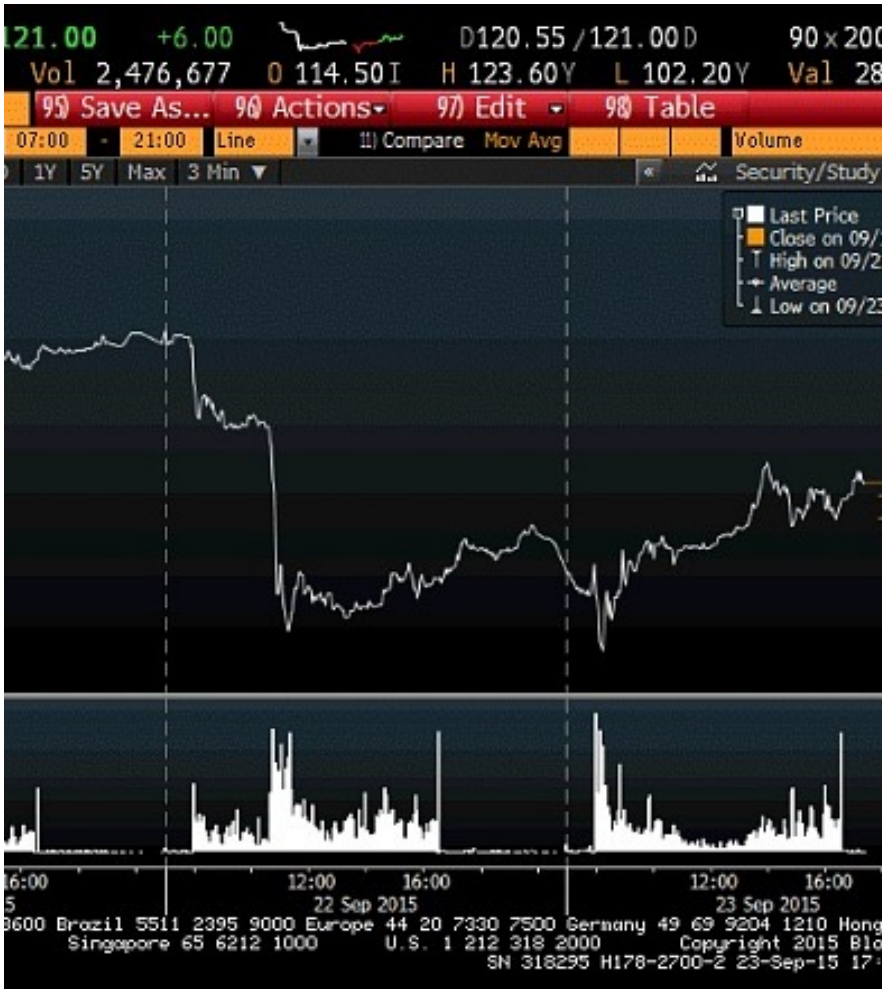

Effect of security threats?

- The possibility of adversarial sample attacks on machine vision systems causes vendors to de-tune
- OK, so someone could project a confusing image on a motorway bridge and cause a crash...
- But someone could already kick some cones into a hole in the road, or drop a brick from a bridge?
- Likely research outcome: better alarms
- Systems involving machine-learning components should be aware when they’re under attack

Intimate partner abuse

- Experienced by 27% of women and 11% of men
- Typically three phases:
  - Physical-control phase where the abuser has access to the survivor's devices
  - Escape phase – finding new home, job etc. Takes on average seven attempts
  - Life-apart phase, where harassment can be an issue; victim may have to change career, or restrict children's online activity in case the abuser turns up
- Much of the standard security advice is unhelpful in such cases – or even plain wrong

Security/safety for abuse survivors

- Standard advice: change your password
- Do you even dare to?
- Could the abuser answer your recovery questions?
- Can you afford to buy another phone, or do you have to plan your escape using a library / work PC?
- How can you use social media if your circle of friends is common with the abuser?
- What should you do as a designer to at least make things no worse for survivors?

Supporting survivors

- Here the enemy is an ‘insider’ in the sense of being the user – some of the time
- Focus on usability at times of high stress, high risk
- Allow users to have multiple accounts
- Someone reviewing your history should not be able to tell if you deleted anything
- BUT you should still be able to capture evidence!
- Push incognito mode, unusual activity notifications, 2-factor authentication (see Consolvo et al paper)

Safety vs security?

- Safety and security have much in common
- They often overlap; they get entangled; they can reinforce each other, and occasionally conflict
- You need to study the context with real care!
- When doing security or safety, it's essential to ask 'safety for whom?' and 'security for whom?' as well as 'safety from what?' and 'security from what?'
- You also need to think methodically about personal and institutional incentives