Personalizing the PC

Prof. Richard Mortier

Systems Research Group,
Department of Computer Science & Technology,
Cambridge University

rmm1002@cam.ac.uk
https://mort.io/
https://mastodon.me.uk/@mort

2026-06-17

The Problem

Then...? 1949: EDSAC!

- We had one computer

© Department of Computer Science and Technology, University of Cambridge. Reproduced by permission.

Then...? 1949: EDSAC!

- We had one computer
- It was expensive

© Department of Computer Science and Technology, University of Cambridge. Reproduced by permission.

Then...? 1949: EDSAC!

- We had one computer
- It was expensive
- It needed many users to keep it busy

Then...? 1949: EDSAC!

- We had one computer
- It was expensive
- It needed many users to keep it busy
- It was designed and built in-house

Then...? 1949: EDSAC!

- We had one computer
- It was expensive
- It needed many users to keep it busy
- It was designed and built in-house
- The first computer seriously used by others than those who built it

The Personal Computer Revolution(s)?

- **Games console** (1970s—), single-purpose consumer product
- **Home computer** (1980s—), general purpose but limited, typically booting into some programming environment
- **PC** (1990s—), general purpose but focused on business uses such as document preparation, spreadsheets, databases
- **Digital assistant** (2000s—), later merging with mobile phones creating the smart phone and tablet

https://www.absolutegeeks.com/article/quick-reads/forgotten-tech-commodore-64/

https://en.wikipedia.org/wiki/File:Atari-2600-Wood-4Sw-Set.jpg

https://www.dosdays.co.uk/computers/IBM%20PC%20(5150)/ibm5150.php

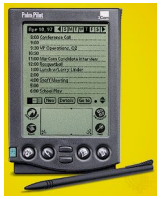

https://www.britannica.com/technology/Palm-Pilot

And Now...? 2026: Oh My!

https://www.apple.com/in/newsroom/2022/10/apple-unveils-completely-redesigned-ipad-in-four-vibrant-colors/

https://www.hippopx.com/en/iphone-apple-phone-cellular-phone-cell-electronics-smartphone-337463

https://www.apple.com/in/newsroom/2022/10/apple-unveils-completely-redesigned-ipad-in-four-vibrant-colors/

https://www.cnet.com/tech/computing/best-laptops-of-ces-2024/

- We each have many computers
  - ...but they don't work together 😔

And Now...? 2026: Oh My!

https://www.apple.com/in/newsroom/2022/10/apple-unveils-completely-redesigned-ipad-in-four-vibrant-colors/

https://www.hippopx.com/en/iphone-apple-phone-cellular-phone-cell-electronics-smartphone-337463

https://www.apple.com/in/newsroom/2022/10/apple-unveils-completely-redesigned-ipad-in-four-vibrant-colors/

https://www.cnet.com/tech/computing/best-laptops-of-ces-2024/

- We each have many computers
  - ...but they don't work together 😔
- They're orders of magnitude cheaper
  - ...comparatively

And Now...? 2026: Oh My!

https://www.apple.com/in/newsroom/2022/10/apple-unveils-completely-redesigned-ipad-in-four-vibrant-colors/

https://www.hippopx.com/en/iphone-apple-phone-cellular-phone-cell-electronics-smartphone-337463

https://www.apple.com/in/newsroom/2022/10/apple-unveils-completely-redesigned-ipad-in-four-vibrant-colors/

https://www.cnet.com/tech/computing/best-laptops-of-ces-2024/

- We each have many computers
  - ...but they don't work together 😔
- They're orders of magnitude cheaper
  - ...comparatively
- And more capable
  - ...not that you can always tell

And Now...? 2026: Oh My!

https://www.apple.com/in/newsroom/2022/10/apple-unveils-completely-redesigned-ipad-in-four-vibrant-colors/

https://www.hippopx.com/en/iphone-apple-phone-cellular-phone-cell-electronics-smartphone-337463

https://www.apple.com/in/newsroom/2022/10/apple-unveils-completely-redesigned-ipad-in-four-vibrant-colors/

https://www.cnet.com/tech/computing/best-laptops-of-ces-2024/

- We each have many computers
  - ...but they don't work together 😔
- They're orders of magnitude cheaper
  - ...comparatively
- And more capable
  - ...not that you can always tell
- We treat them as consumer goods
  - ...very few people want to become distributed system experts

So What?

Does this matter? I argue that **yes, it does**:

- Data management is a mess:
  - ...files, manually and inconsistently synced and versioned
  - ...by the (usually non-expert) user
- Reliance on cloud services in the face of:
  - ...intermittent network connectivity
  - ...lack of meaningful privacy controls
- Identity management is mis-targeted:
  - ...identifying users not applications
  - ...applications ineffectively sandboxed

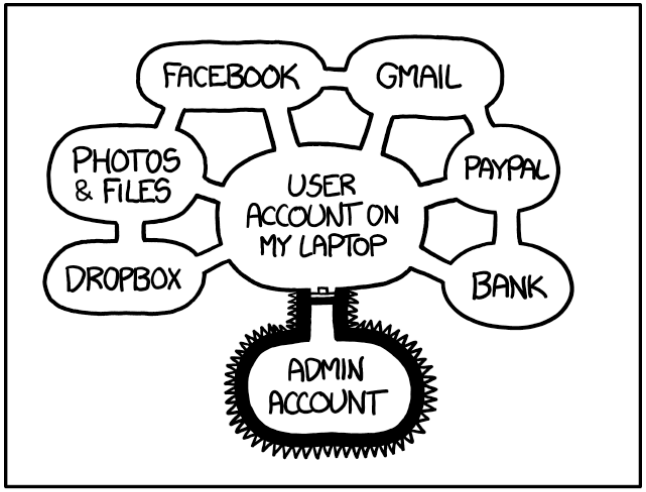

https://xkcd.com/1200/

The Vision

What is a “Personal Computer” *today*?

These are all individual devices; many are mobile!

- Laptop. Tablet. Phone. Watch. Home PC.
- All have considerable compute, memory, storage, networking
  - Though some have more than others (GB vs TB; Mb/s vs Gb/s)
- No operating system targets the core use-case:
  - Linux distributions build for developers, or ape (old) Windows
  - Apple builds for a closed ecosystem with iCloud at the centre
  - Windows is ad centric with Microsoft M365 at the centre
  - Android is device centric with Google at the centre

What is a “Personal Computer” *today*?

These are all individual devices; many are mobile!

- Laptop. Tablet. Phone. Watch. Home PC.
- All have considerable compute, memory, storage, networking
  - Though some have more than others (GB vs TB; Mb/s vs Gb/s)
- No operating system targets the core use-case:
  - Linux distributions build for developers, or ape (old) Windows
  - Apple builds for a closed ecosystem with iCloud at the centre
  - Windows is ad centric with Microsoft M365 at the centre
  - Android is device centric with Google at the centre

Claim #1 — a modern “personal computer” is a small **personal constellation** comprising several discrete devices

What is a “Personal Computer” *today*?

Claim #1 — a modern “personal computer” is a small **personal constellation** comprising several discrete devices

From Mainframes to Minicomputers to … what?

- What about the system software they run?
- Fundamental design assumptions haven’t evolved for decades:
  - Multiple users to keep an expensive machine busy
  - Running applications without really defining what one is
  - Using dynamic shared libraries to save space at the cost of package management complexity

What is a “Personal Computer” *today*?

Claim #1 — a modern “personal computer” is a small **personal constellation** comprising several discrete devices

Claim #2 — modern system software is still built assuming one multi-user single-CPU resource-constrained computer

What is a “Personal Computer” *today*?

Claim #1 — a modern “personal computer” is a small **personal constellation** comprising several discrete devices

Claim #2 — modern system software is still built assuming one multi-user single-CPU resource-constrained computer

It’s really about the data interaction **not the** device management

- As Xkcd has it “*If someone steals my laptop while I’m logged in, they can read my email, take my money, and impersonate me to my friends, but at least they can’t install drivers without my permission.*”
- We share data and consume services across multiple devices

What is a “Personal Computer” *today*?

Claim #1 — a modern “personal computer” is a small **personal constellation** comprising several discrete devices

Claim #2 — modern system software is still built assuming one multi-user single-CPU resource-constrained computer

Claim #3 — need to shift the focus to supporting users to manage interactions around data

What is a “Personal Computer” *today*?

Claim #1 — a modern “personal computer” is a small **personal constellation** comprising several discrete devices

Claim #2 — modern system software is still built assuming one multi-user single-CPU resource-constrained computer

Claim #3 — need to shift the focus to supporting users to manage interactions around data

- Aligns with pillars of **Human-Data Interaction**:
  - **Legibility**, seeing and understanding data collected about you
  - **Agency**, controlling and influencing data collection about you
  - **Negotiability**, collaborating and negotiating with others

How would a Personal Constellation be better?

- You're playing music from your 2TB iPad on a bluetooth speaker at home with your 64GB phone in hand
  - Why can't you see the current track progress and playlist on your phone?
  - Why can't you control volume and skip track from your phone?

How would a Personal Constellation be better?

- You're playing music from your 2TB iPad on a bluetooth speaker at home with your 64GB phone in hand

- You leave the house to get in your car carrying your phone and your laptop
  - Why doesn't playback continue from your phone on your car's bluetooth speakers?
  - Why did you have to manually select the tracks to synchronise to your phone from your iPad?
  - Why did any of the above involve cloud-hosted services?

How would a Personal Constellation be better?

- You're playing music from your 2TB iPad on a bluetooth speaker at home with your 64GB phone in hand

- You leave the house to get in your car carrying just your phone

- You drive to meet a colleague and wish to share data with them
  - Why can't you airdrop a URL from a browser on your laptop by bumping phones?
  - Why can't you let them play one of your music tracks while you're together, without affecting your own playback metadata?

Inspirational Interlude

Homework

- Revisited a key design assumption:
  - Networks have expert management
- Used ethnographic techniques to explore people's understanding and use of their home networks
  - E.g., Network "walk-throughs": users documenting their home networks
- Lead to looking again at underlying platform technology
  - Relatively minor modifications lead to much greater flexibility

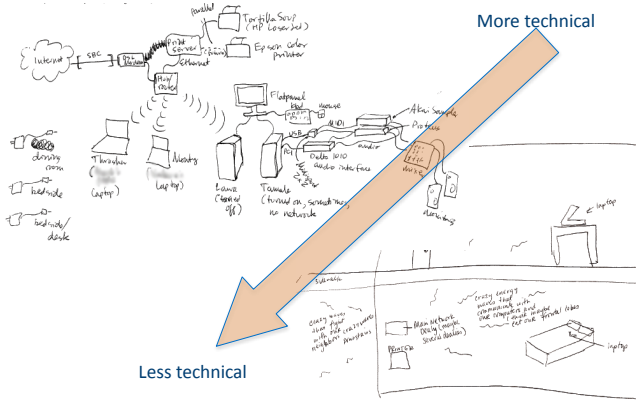

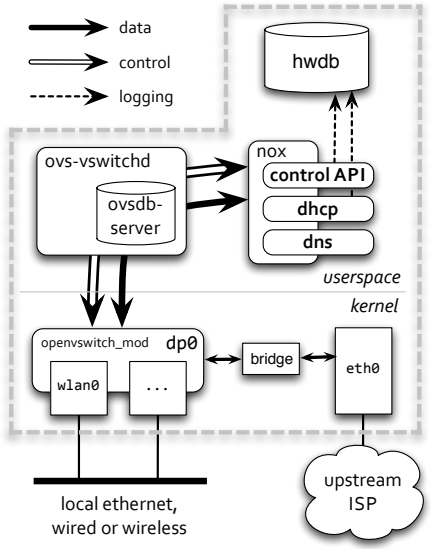

Homework

- Users control their networks more effectively once the technology is made legible
  - "*MultiNet: reducing interaction overhead in domestic wireless networks*", Brown et al, ACM CHI 2013
- Users constructed policy rules that exercised nuanced control over use of their networks
  - "*House Rules: the collaborative nature of policy in domestic networks*", Crabtree et al, PUC 2015

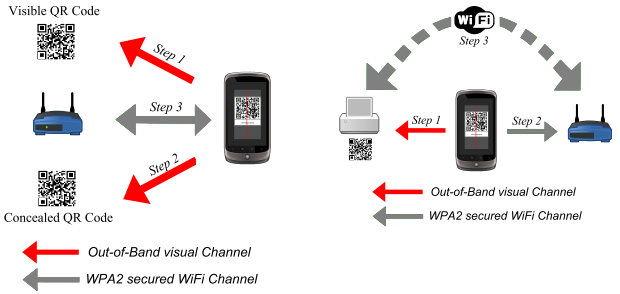

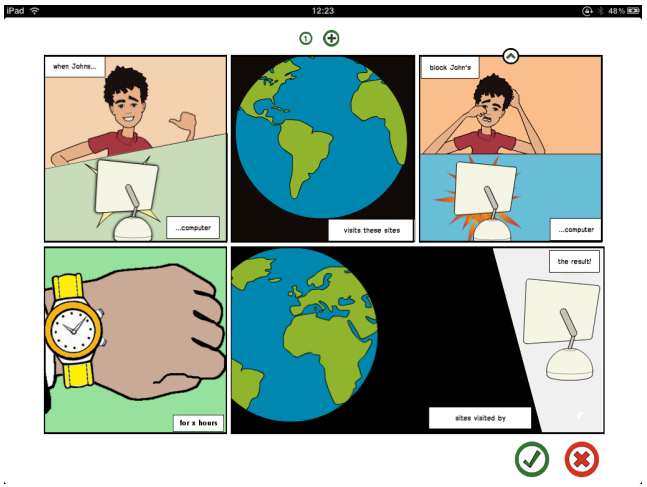

UCN: Personal Information Hub

User-Centric Networking

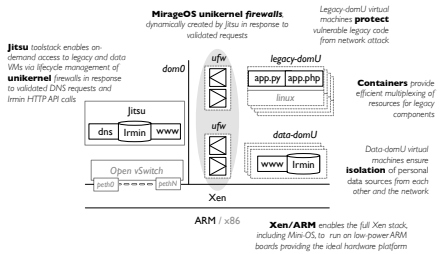

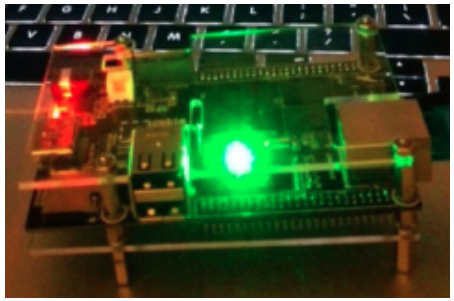

- Mirage unikernels used to create a Personal Information Hub (PIH)
  - "*Unikernels: library operating systems for the cloud*", Madhavapeddy et al, ASPLOS 2013
- Data in distinct stores, access controlled via network permissions
- Authentication creates a lightweight user-firewall configured to limit and audit access to data store
  - "*Jitsu: Just-In-Time Summoning of Unikernels*", USENIX NSDI 2015

Databox

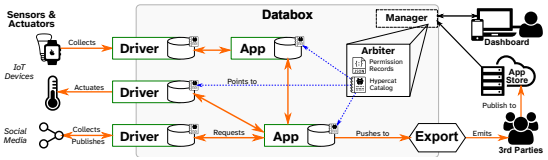

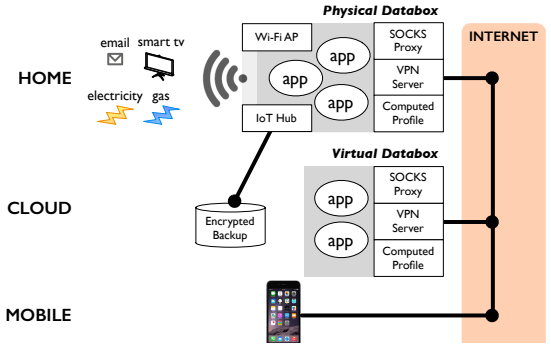

- Built on ideas from UCN PIH
- Move the code to the data, not the data to the code
  - Host data processing code as containerised apps
  - Supported by manifests, drivers, stores expressing and enforcing intents
- Allows constrained delegation of access to data with audit
  - "*Personal Data Management with the Databox: What's Inside the Box?*", Mortier et al, ACM CAN 2016

Databox: An observation

> *Current systems assume a root identity with multiple subordinate user identities – but we have one user running multiple applications accessing shared data "...privacy dissolves into ...relationship management practices ...'privacy' has little utility as a focus for design ...a more productive way forward would be to concentrate on supporting people's evident interest in managing their relationships in and with the networked world.*
>
> *— Crabtree, Tolmie, Knight (ACM CSCW 2017)*[1]

[1] https://doi.org/10.1007/s10606-017-9276-y

So What?

- **One** user, **multiple** applications, **shared** data
- Relationship management between user and
  - their multiple **devices**,
  - their multiple **data silos**,
  - and multiple **other users**

So What?

- **One** user, **multiple** applications, **shared** data
- Relationship management between user and
  - their multiple **devices**,
  - their multiple **data silos**,
  - and multiple **other users**
- How do we ensure the right data is on the right device at the right time?
  - The right device **might not be my device**: we transiently share and collaborate
  - The right time **has a start and an end**: we need to be able to revoke access

So What?

- **One** user, **multiple** applications, **shared** data
- Relationship management between user and
  - their multiple **devices**,
  - their multiple **data silos**,
  - and multiple **other users**
- How do we ensure the right data is on the right device at the right time?
- The user understands their relationships
  - Don't guess!
  - Give them the controls so they can decide what to do
    - ...with sensible defaults

Scarecrow Solutions

So What?

https://editorial01.shutterstock.com/preview-440/204689c/d6e8b6ee/Shutterstock_204689c.jpg

- What underlying abstractions we would come up with if we started again with a clean-slate?
  - Single-user? Distributed by default? Data-focused? Dynamic linking?

Homework user-centric infrastructure

UCN focus on service management

Databox we need a fundamentally different approach

So What?

https://editorial01.shutterstock.com/preview-440/204689c/d6e8b6ee/Shutterstock_204689c.jpg

- What underlying abstractions we would come up with if we started again with a clean-slate?
  - Single-user? Distributed by default? Data-focused? Dynamic linking?

Homework user-centric infrastructure

UCN focus on service management

Databox we need a fundamentally different approach

Just what is an operating system for?

Storage

- File-systems that are robust to failure
- Backup as a built-in characteristic
- *Draw on*: embedded systems work?
- *Means that*: **your constellation makes your data resilient**
  - ...to device failure, intermittent connectivity, storage removal

Storage

- File-systems that are robust to failure
- Backup as a built-in characteristic
- *Draw on*: embedded systems work?
- *Means that*: **your constellation makes your data resilient**
  - ...to device failure, intermittent connectivity, storage removal

- Distributed by nature
- Automatically managed, manually hinted
- *Draw on*: CRDTs? statistics for predicting access patterns?
- *Means that*: **your data is as locally available as possible**

Networking

- Securely and uniformly interconnected, on-demand
- Policy-enforcement according to stated behaviour
- *Draw on*: per-application flow monitoring? MUD-style policy manifests?
- *Means that*: **your devices always see each other, securely**

Networking

- Securely and uniformly interconnected, on-demand
- Policy-enforcement according to stated behaviour
- *Draw on*: per-application flow monitoring? MUD-style policy manifests?
- *Means that*: **your devices always see each other, securely**

- Intersect application policies with user's
- Automatically determined and managed
- *Draw on*: statistics for anomaly detection? policy inference? Homework policy UX?
- *Means that*: **you can detect and prevent misbehaviour**

Applications

- Well-defined, self-contained, principal of resource management
- Distinguish service from UI
- *Draw on*: unikernels? deduplication?
- *Means that*:
  - **you can compute locally and render conveniently**
  - **you can see and manage what each application is doing**

Applications

- Well-defined, self-contained, principal of resource management
- Distinguish service from UI
- *Draw on*: unikernels? deduplication?
- *Means that*:
  - **you can compute locally and render conveniently**
  - **you can see and manage what each application is doing**

- Centrally managed build and deployment
- Automatically-upated, statically-linked, richly described
- *Draw on*: TUF? in-toto? Cosmopolitan? Flatpak, Snap, 0install &c
- *Means that*: **your constellation acts as one distributed device**

Data & Devices

- Automatic real-time synchronisation across devices (video)
- Specialisation of storage formats
- *Draw on*: statistics for optimization? CRDTs? novel consensus systems?
- *Means that*: **intermittent connectivity is handled seamlessly**

Data & Devices

- Automatic real-time synchronisation across devices (video)
- Specialisation of storage formats
- *Draw on*: statistics for optimization? CRDTs? novel consensus systems?
- *Means that*: **intermittent connectivity is handled seamlessly**

- Explicit configuration, system lockfiles, long-term distribution
- Focus on containment and predictability
- *Draw on*: NixOS? Securix? QubesOS?
- *Means that*: **system management and upgrade becomes trivial**

Conclusions: What to do?

- User studies
  - What do people really want? What do they really do?
- Resource tracking and identity management
  - What do current OSs really do? Is it enough?
  - What structurally changes when AppID replaces UID/GID?
  - Can eBPF/ETW/dtrace support pervasive data provenance?
- Application definition
  - What does it look like to integrate CRDTs with storage?
  - How far can polyglot binaries and fancy linking be pushed?
  - What does the development process look like when supporting the application compute/render split?

Questions?

Feedback!

Homework, https://homework.github.io/

- The platform, https://doi.org/10.1145/2380116.2380143
- A design/sociological take on it, https://doi.org/10.1145/2317956.2318039
- A better way to securely associated with wireless access points, https://doi.org/10.1145/2470654.2466208
- An intuitive comic strip policy interface, https://doi.org/10.1007/s00779-014-0771-6

Databox, https://github.com/me-box/databox?tab=readme-ov-file#published-work

- Personal data: thinking inside the box, https://doi.org/10.7146/aahcc.v1i1.21312
- Personal Data Management with the Databox: What's inside the Box? https://doi.org/10.1145/3010079.3010082
- Repacking 'Privacy' for a Networked World, https://doi.org/10.1007/s10606-017-9276-y
- The practical politics of sharing personal data, https://doi.org/10.1007/s00779-017-1071-8
- Building accountability into the Internet-of-Things, https://doi.org/10.1007/s40860-018-0054-5
- Occupancy-as-a-(Private-)Service, https://doi.org/10.1145/3277893.3277894
- Privacy-preserving personal model training, https://doi.org/10.1109/IoTDI.2018.00024

Consensus at the Edge

- Rearchitecting Kubernetes for the Edge, https://doi.org/10.1145/3434770.3459730
- Paxos vs Raft: Have we reached consensus on distributed consensus?, https://doi.org/10.1145/3380787.3393681
- Examining Raft's behaviour during partial network failures, https://doi.org/10.1145/3447851.3458739