

Automerge:

A new foundation for
collaboration software

MARTIN KLEPPMANN
UNIVERSITY OF CAMBRIDGE

@martinkl



UNIVERSITY OF
CAMBRIDGE

Ink & Switch

Thank you to my supporters

LEVERHULME
TRUST



NOKIA Bell Labs

Crowdfunding supporters include: Ably, Adrià Arcarons, Chet Corcos, Macrometa, Mintter, David Pollak, Prisma, RelationalAI, SoftwareMill, Talent Formation Network, Adam Wiggins

<https://www.patreon.com/martinkl>

Example: Text editing



Example: Text editing

insert "World"
after "Hello"



"Hello!"

"Hello World!"

time


insert ":-)"
after "!"



"Hello!"

"Hello! :-)"

Example: Text editing



hack on code

time







hack on code

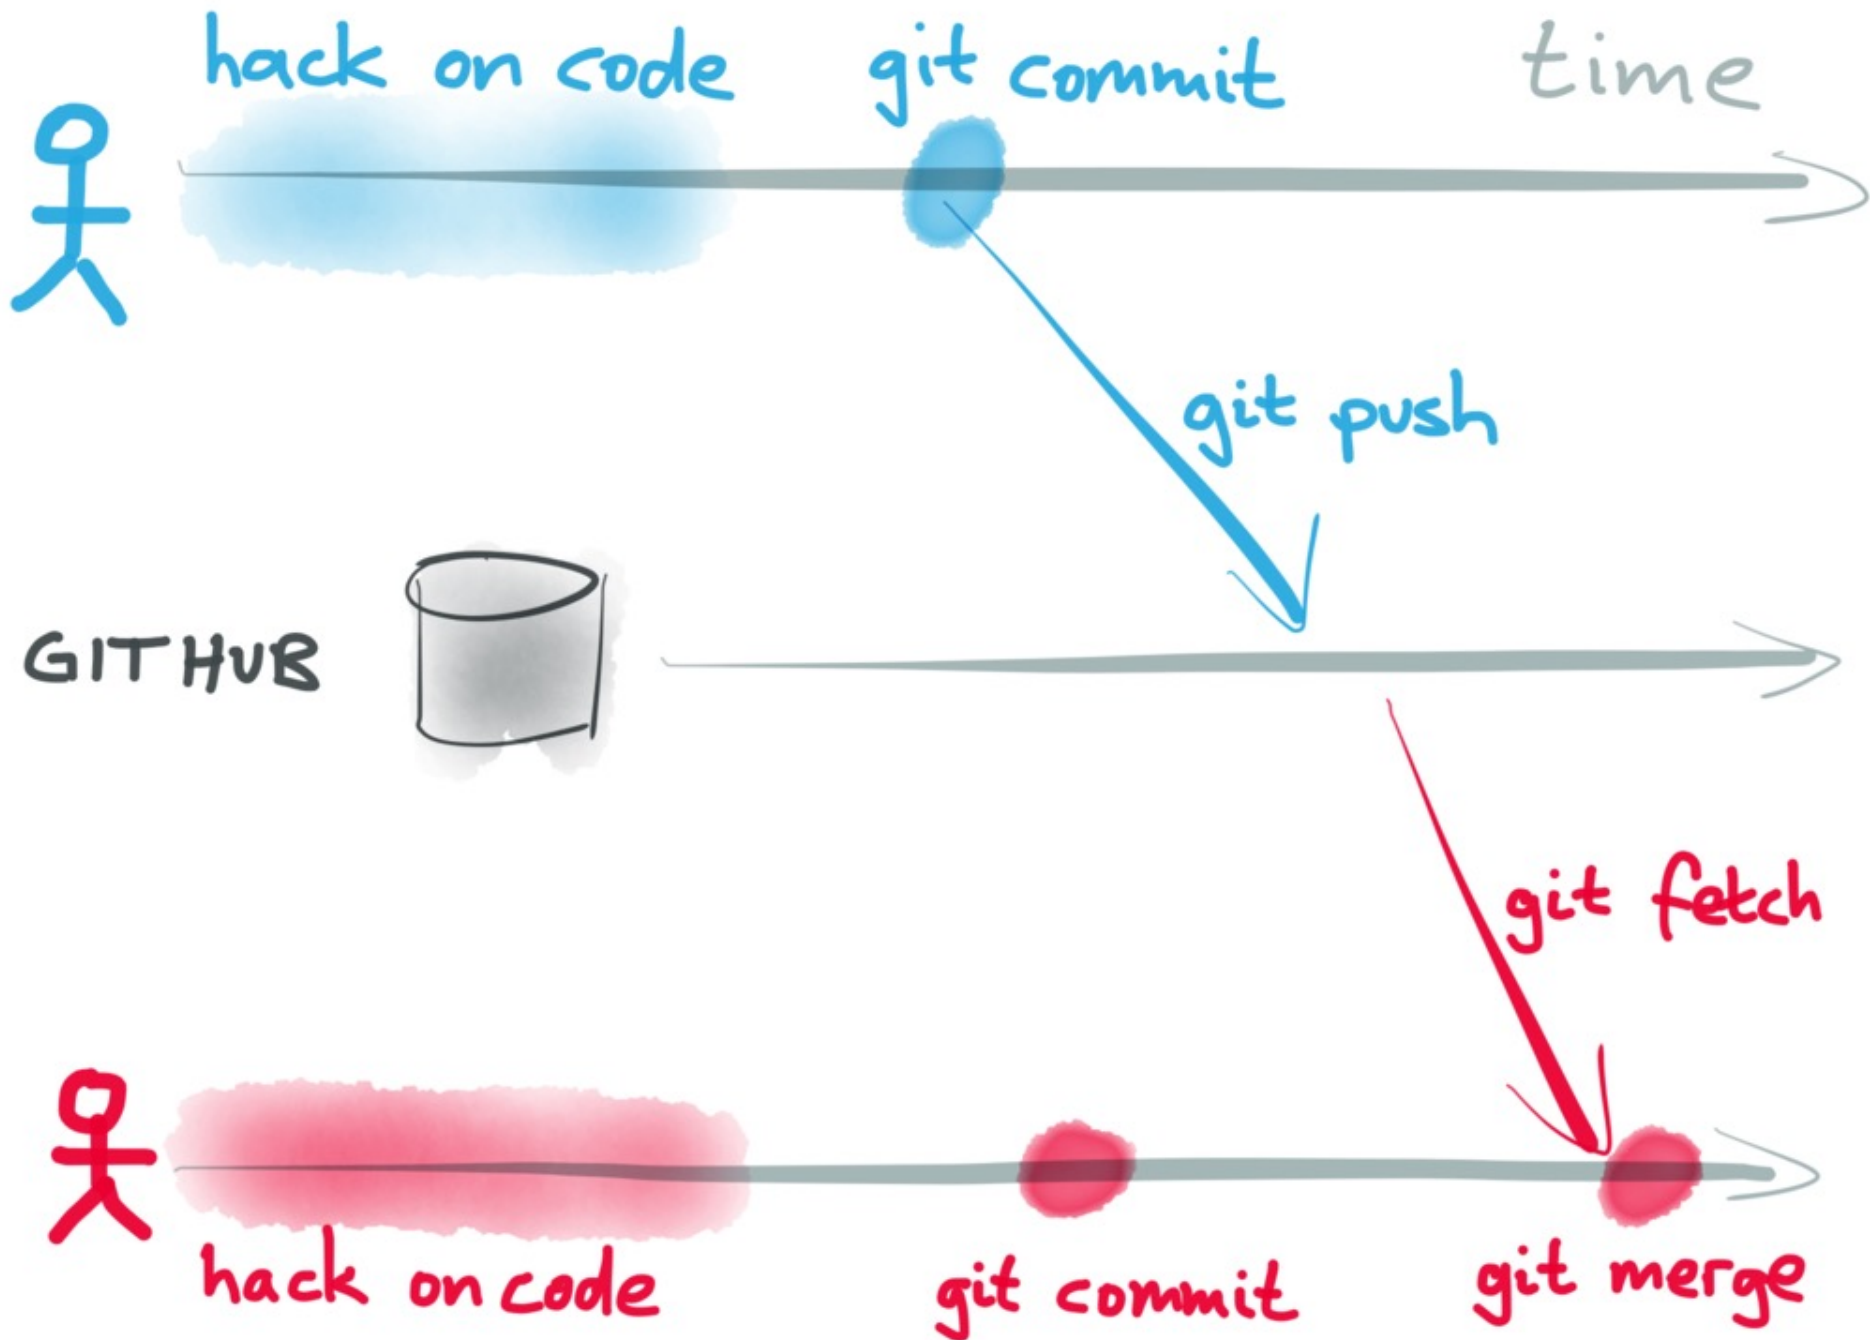
git commit

time









COLLABORATIVE APPLICATIONS



Google Docs



Google Sheets



Office 365

Overleaf

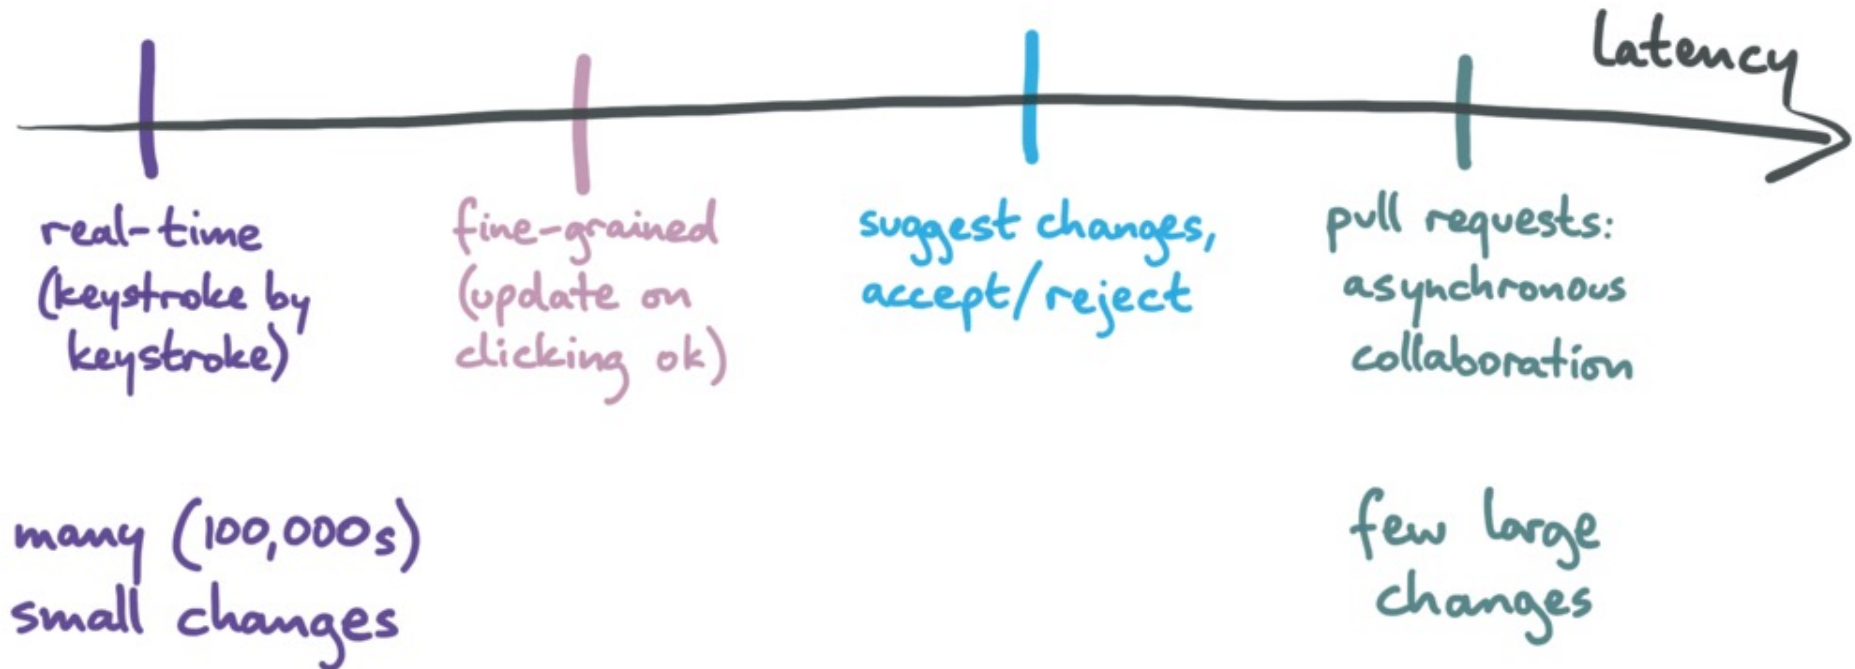
Trello



Figma

Collaboration

Syncing changes between several users



AUTOMERGE: Branching and merging

```
{ "todos": [  
  { "title": "buy milk",  
    "done": false },  
  { "title": "water plants",  
    "done": false }  
]}
```


AUTOMERGE: Branching and merging

Automerge.change

USER A:

```
{ "todos": [  
  { "title": "buy milk",  
    "done": false },  
  { "title": "water plants",  
    "done": true }  
]}
```

```
{ "todos": [  
  { "title": "buy milk",  
    "done": false },  
  { "title": "water plants",  
    "done": false }  
}]}
```

AUTOMERGE: Branching and merging

Automerge.change

USER A:

```
{ "todos": [  
  { "title": "buy milk",  
    "done": false },  
  { "title": "water plants",  
    "done": true }  
]}
```


```
{ "todos": [  
  { "title": "buy milk",  
    "done": false },  
  { "title": "water plants",  
    "done": false }  
]}
```


Automerge.change

USER B:

```
{ "todos": [  
  { "title": "buy milk",  
    "done": false },  
  { "title": "water plants",  
    "done": false },  
  { "title": "do laundry",  
    "done": false }  
]}
```


AUTOMERGE: Branching and merging





Automerge

<https://github.com/automerge/automerge>

AUTOMERGE: "Git for your app's data"

```
{ "todos": [
```

```
  { "title": "buy milk", "done": false },
```

```
  { "title": "water plants", "done": false }
```

```
]
```

AUTOMERGE: "Git for your app's data"

```
{ "todos": [  
  { "title": "buy milk", "done": false },  
  { "title": "water plants", "done": false }  
]}
```

```
after = Automerge.change(before, "add new item", doc => {  
  doc.todos.push({title: "do laundry", done: false});  
});
```

AUTOMERGE: "Git for your app's data"

```
{ "todos": [  
  { "title": "buy milk", "done": false },  
  { "title": "water plants", "done": false },  
  { "title": "do laundry", "done": false } ← added  
]}
```

```
after = Automerge.change(before, "add new item", doc => {  
  doc.todos.push({ title: "do laundry", done: false });  
});
```


AUTOMERGE: "Git for your app's data"

```
{ "todos": [  
  { "title": "buy milk", "done": false },  
  { "title": "water plants", "done": false },  
  { "title": "do laundry", "done": false } ← added  
]}
```

```
after = Automerge.change(before, "add new item", doc => {  
  doc.todos.push({ title: "do laundry", done: false });  
});
```

← append item to list

```
Automerge.change(state, doc => {  
  doc.todos.push({title: "do laundry", done: false});  
});
```

```
Automerge.change(state, doc => {  
  doc.todos.push({title: "do laundry", done: false});  
});
```

operation
log

```
{op: "makeMap", id: "5a", obj: "1a", elemID: "2a", insert: true}  
{op: "assign", id: "6a", obj: "5a", key: "title", value: "do laundry"}  
{op: "assign", id: "7a", obj: "5a", key: "done", value: false}
```



```
Automerge.change(state, doc => {  
  doc.todos.push({title: "do laundry", done: false});  
});
```

operation
log

```
{op: "makeMap", id: "5a", obj: "1a", elemID: "2a", insert: true}  
{op: "assign", id: "6a", obj: "5a", key: "title", value: "do laundry"}  
{op: "assign", id: "7a", obj: "5a", key: "done", value: false}
```

compressed
binary
encoding

```
Uint8Array([0x85, 0x6f, 0x4a, 0x83, ...])
```

```
Automerge.change(state, doc => {  
  doc.todos.push({title: "do laundry", done: false});  
});
```

operation
log

```
{op: "makeMap", id: "5a", obj: "1a", elemID: "2a", insert: true}  
{op: "assign", id: "6a", obj: "5a", key: "title", value: "do laundry"}  
{op: "assign", id: "7a", obj: "5a", key: "done", value: false}
```

compressed
binary
encoding

```
Uint8Array([0x85, 0x6f, 0x4a, 0x83, ...])
```

write to disk, send over network

{title: "Water plants", done: false}

{op: "makeMap", id: "1a"}

{op: "assign", id: "2a", obj: "1a", key: "title", value: "Water plants", overwrites: []}

{op: "assign", id: "3a", obj: "1a", key: "done", value: false, overwrites: []}

`{title: "Water plants", done: false}`

`{op: "makeMap", id: "1a"}`

`{op: "assign", id: "2a", obj: "1a", key: "title", value: "Water plants", overwrites: []}`

`{op: "assign", id: "3a", obj: "1a", key: "done", value: false, overwrites: []}`

`doc.todos[0].done = true`

`{op: "assign", id: "4a", obj: "1a", key: "done", value: true, overwrites: ["3a"]}`

{title: "Water plants", done: false}

{op: "makeMap", id: "1a"}

{op: "assign", id: "2a", obj: "1a", key: "title", value: "Water plants", overwrites: []}

{op: "assign", id: "3a", obj: "1a", key: "done", value: false, overwrites: []}

doc.todos[0].done = true

{op: "assign", id: "4a", obj: "1a", key: "done", value: true, overwrites: ["3a"]}

{title: "Water plants", done: false}

{op: "makeMap", id: "1a"}

{op: "assign", id: "2a", obj: "1a", key: "title", value: "Water plants", overwrites: []}

{op: "assign", id: "3a", obj: "1a", key: "done", value: false, overwrites: []}

doc.todos[0].done = true

{op: "assign", id: "4a", obj: "1a", key: "done", value: true, overwrites: ["3a"]}

MANUAL CONFLICT RESOLUTION

```
doc.todos[0].deadline  
= "2021-07-10"
```

```
doc.todos[0].deadline  
= "2021-07-14"
```

merge

A diagram illustrating a merge operation. Two rectangular boxes, one above the other, contain code snippets. The top box has 'doc.todos[0].deadline = "2021-07-10"' and the bottom box has 'doc.todos[0].deadline = "2021-07-14"'. To the right of the boxes, the word 'merge' is written. Two curved arrows originate from the right side of each box and point towards a single arrowhead on the right, indicating the merging of the two states.

`doc.todos[0].deadline = "2021-07-10"`

`{op: "assign", id: "5a", obj: "1a", key: "deadline", value: "2021-07-10",
overwrites: []}`

`doc.todos[0].deadline = "2021-07-14"`

`{op: "assign", id: "5b", obj: "1a", key: "deadline", value: "2021-07-14",
overwrites: []}`

`doc.todos[0].deadline = "2021-07-10"`

`{op: "assign", id: "5a", obj: "1a", key: "deadline", value: "2021-07-10",
overwrites: []}`

`doc.todos[0].deadline = "2021-07-14"`

`{op: "assign", id: "5b", obj: "1a", key: "deadline", value: "2021-07-14",
overwrites: []}`

MANUAL CONFLICT RESOLUTION

`doc.todos[0].deadline`
`= "2021-07-10"`

`doc.todos[0].deadline`
`= "2021-07-14"`

merge

Automerge.
`getConflicts(doc.todos[0],`
`"deadline")`

`= {`
 `5a: "2021-07-10",`
 `5b: "2021-07-14"`
 `}`

TIME TRAVEL

Automerge.getHistory(*state*)

TIME TRAVEL

Automerge.getHistory(*state*)

⇒ [{ change: { message: "Add todo item", ... },
 snapshot: { todos: [{ title: "Buy milk", ... }, ...] } },
 { change: { message: "Mark item as done", ... },
 snapshot: { todos: [{ title: "Buy milk", ... }, ...] } },
 ...
]

{op: "makeMap", id: "1a"}

{op: "assign", id: "2a", obj: "1a", key: "title", value: "Water plants", overwrites: []}

{op: "assign", id: "3a", obj: "1a", key: "done", value: false, overwrites: []}

{op: "assign", id: "4a", obj: "1a", key: "done", value: true, overwrites: ["3a"]}

{op: "makeMap", id: "1a"}

{op: "assign", id: "2a", obj: "1a", key: "title", value: "Water plants", overwrites: []}

{op: "assign", id: "3a", obj: "1a", key: "done", value: false, overwrites: []}

{op: "assign", id: "4a", obj: "1a", key: "done", value: true, overwrites: ["3a"]}

convert into table of ops

op	id	obj	key	value	overwritten by
makeMap	1a	root	todo		{}
assign	3a	1a	done	false	{4a}
assign	4a	1a	done	true	{}
assign	2a	1a	title	"Water plants"	{}

Sort order: 1. by object ID
2. if map: lexicographic by key
2. if list: order of list elements
3. by operation ID

op	id	obj	key	value	overwritten by
makeMap	1a	root	todo		$\{\}$
assign	3a	1a	done	false	$\{4a\}$
assign	4a	1a	done	true	$\{\}$
assign	2a	1a	title	"Water plants"	$\{\}$

Identify document version by version vector
 e.g. $V = \{a: 3, b: 4\}$

op	id	obj	key	value	overwritten by
makeMap	1a	root	todo		$\{\}$
assign	3a	1a	done	false	$\{4a\}$
assign	4a	1a	done	true	$\{\}$
assign	2a	1a	title	"Water plants"	$\{\}$

Identify document version by version vector
 e.g. $V = \{a: 3, b: 4\}$

Visibility rule: operation with $ID = (ctr_{10}, node_{10})$ and
 $overwrittenBy = \{(ctr_1, node_1), (ctr_2, node_2), \dots\}$ is visible at document
 version V iff $ctr_{10} \leq V[node_{10}]$ and

$\nexists (ctr_i, node_i) \in overwrittenBy. ctr_i \leq V[node_i].$

Identify document version by version vector
e.g. $V = \{a: 3, b: 4\}$

Visibility rule: operation with $ID = (ctr_{i0}, node_{i0})$ and
 $overwrittenBy = \{(ctr_1, node_1), (ctr_2, node_2), \dots\}$ is visible at document
version V iff $ctr_{i0} \leq V[node_{i0}]$ and
 $\nexists (ctr_i, node_i) \in overwrittenBy. ctr_i \leq V[node_i]$.

Like MVCC in databases with snapshot isolation!

Identify document version by version vector
e.g. $V = \{a: 3, b: 4\}$

Visibility rule: operation with $ID = (ctr_{id}, node_{id})$ and
 $overwrittenBy = \{(ctr_1, node_1), (ctr_2, node_2), \dots\}$ is visible at document
version V iff $ctr_{id} \leq V[node_{id}]$ and
 $\nexists (ctr_i, node_i) \in overwrittenBy. ctr_i \leq V[node_i]$.

Like MVCC in databases with snapshot isolation!

For a given key / list element:

- no ops visible \Rightarrow deleted
- one op visible \Rightarrow current value
- multiple ops visible \Rightarrow conflict (concurrent assignment)

COLUMNAR ENCODING (simplified)

operation ID		reference element ID		inserted character		deleted by opID	
counter	actor	counter	actor	length	UTF-8	counter	actor
1	A	—	—	1	"H"	—	—
2	A	1	A	1	"e"	—	—
3	A	2	A	1	"l"	—	—
4	A	3	A	1	"l"	—	—
5	A	4	A	1	"l"	7	B
6	A	5	A	1	"o"	—	—

COLUMNAR ENCODING (simplified)

operation ID		reference element ID		inserted character		deleted by opID	
counter	actor	counter	actor	length	UTF-8	counter	actor
1	A	—	—	1	"H"	—	—
2	A	1	A	1	"e"	—	—
3	A	2	A	1	"l"	—	—
4	A	3	A	1	"l"	—	—
5	A	4	A	1	"l"	7	B
6	A	5	A	1	"o"	—	—

→ 1, 2, 3, 4, 5, 6

delta-encode to 1, 1, 1, 1, 1, 1

run-length encode to (6, 1)

LEB128 encodes this in 2 bytes

COLUMNAR ENCODING (simplified)

operation ID		reference element ID		inserted character		deleted by opID	
counter	actor	counter	actor	length	UTF-8	counter	actor
1	A	—	—	1	"H"	—	—
2	A	1	A	1	"e"	—	—
3	A	2	A	1	"l"	—	—
4	A	3	A	1	"l"	—	—
5	A	4	A	1	"l"	7	B
6	A	5	A	1	"o"	—	—

→ make a lookup table: $\{"A": 0, "B": 1\}$
→ 0, 0, 0, 0, 0, 0
→ run-length encode to (6, 0)
→ LEB128 encodes in 2 bytes

COLUMNAR ENCODING (simplified)

operation ID		reference element ID		inserted character		deleted by opID	
counter	actor	counter	actor	length	UTF-8	counter	actor
1	A	—	—	1	"H"	—	—
2	A	1	A	1	"e"	—	—
3	A	2	A	1	"l"	—	—
4	A	3	A	1	"l"	—	—
5	A	4	A	1	"l"	7	B
6	A	5	A	1	"o"	—	—

just concatenate the UTF-8
byte sequences → "Hello" (6 bytes)
(use length column to separate again)

COLUMNAR ENCODING (simplified)

operation ID		reference element ID		inserted character		deleted by opID	
counter	actor	counter	actor	length	UTF-8	counter	actor
1	A	—	—	1	"H"	—	—
2	A	1	A	1	"e"	—	—
3	A	2	A	1	"l"	—	—
4	A	3	A	1	"l"	—	—
5	A	4	A	1	"l"	7	B
6	A	5	A	1	"o"	—	—

Plus some additional metadata (e.g. timestamp and range of opID counter values for each change)
⇒ can reconstruct any past document state


Automerge compression benchmark

Benchmark data: keystroke-by-keystroke editing trace of a text file (LaTeX source of a research paper) containing 182,315 single-character insertions and 77,463 single-character deletions, timestamped with 1-second granularity.

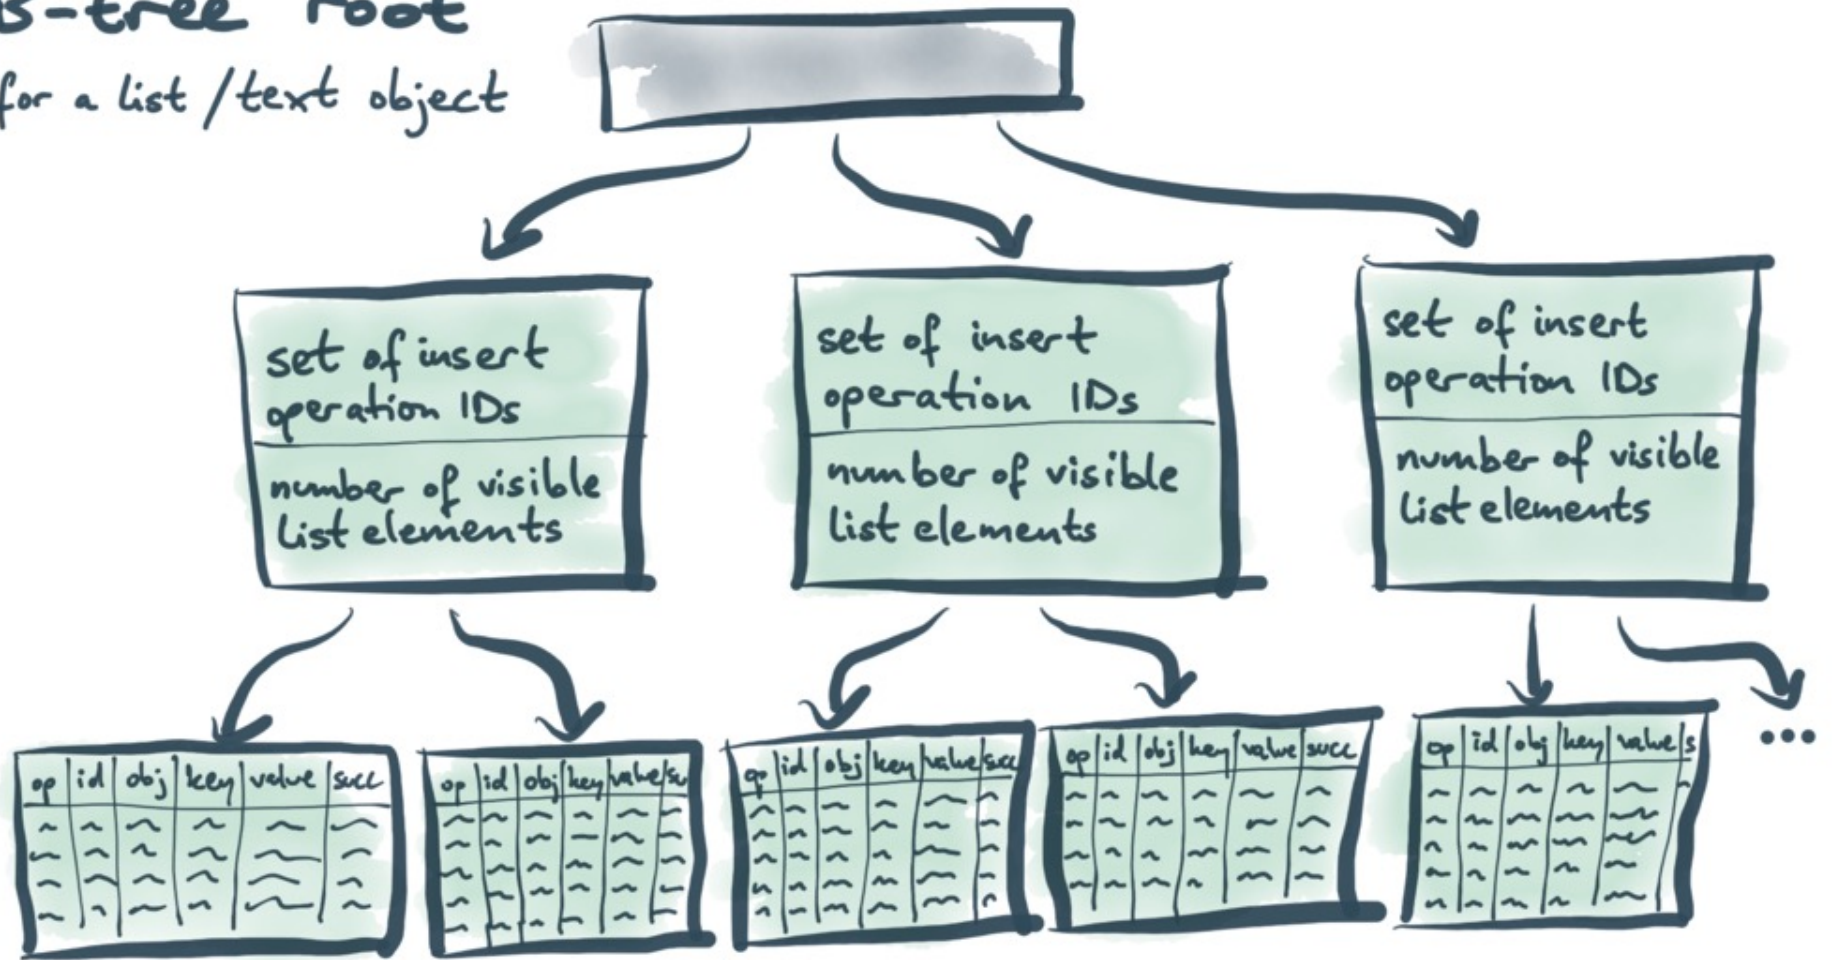
As individual changes: 33.7 MB (130 bytes/operation)

As compressed document with full edit history: 184 kB (0.7 bytes/operation)

Breakdown of compressed columnar file contents



B-tree root
for a list/text object



Local insert: find insertion position by index, counting only visible elements, to translate index into opID of last insert at that index


Remote insert: find subtree containing position ID, add op to B-tree, compute index of inserted element based on number of preceding visible elements

IMPLEMENTATION

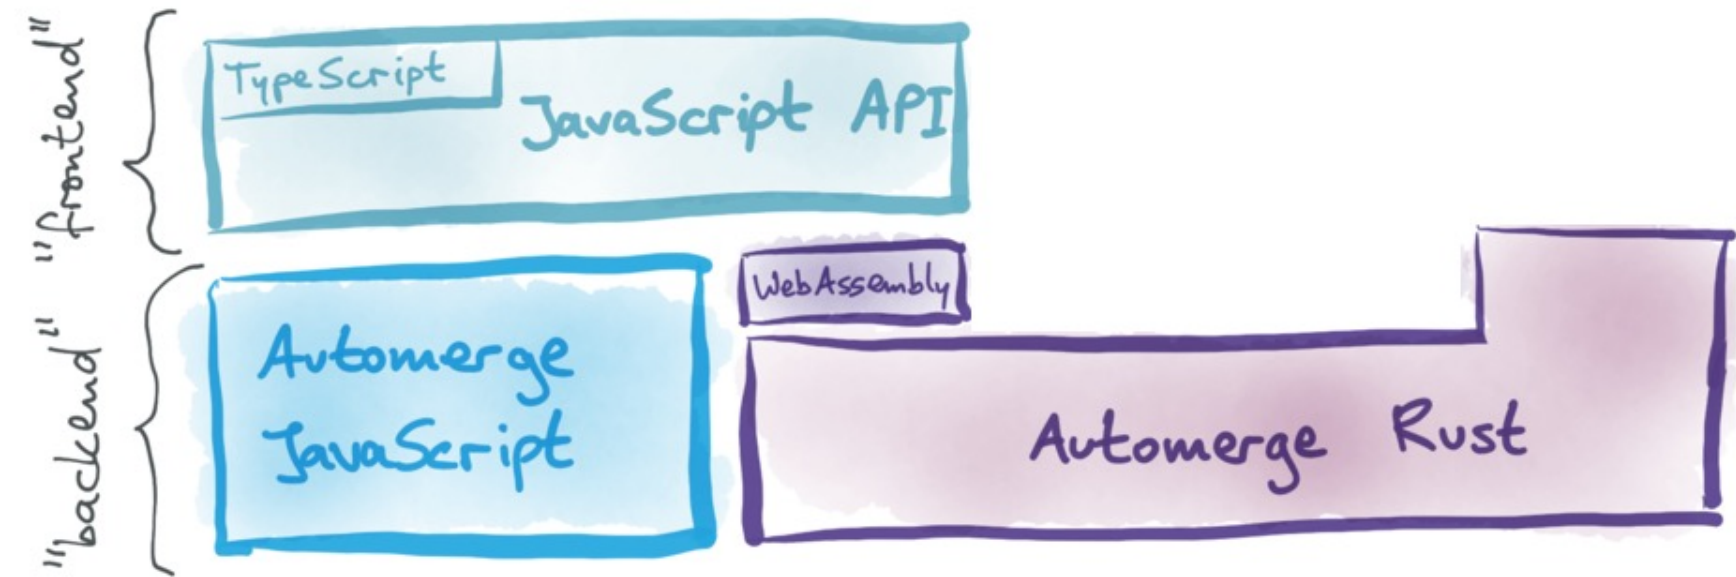
TypeScript
JavaScript API

Automerge
JavaScript

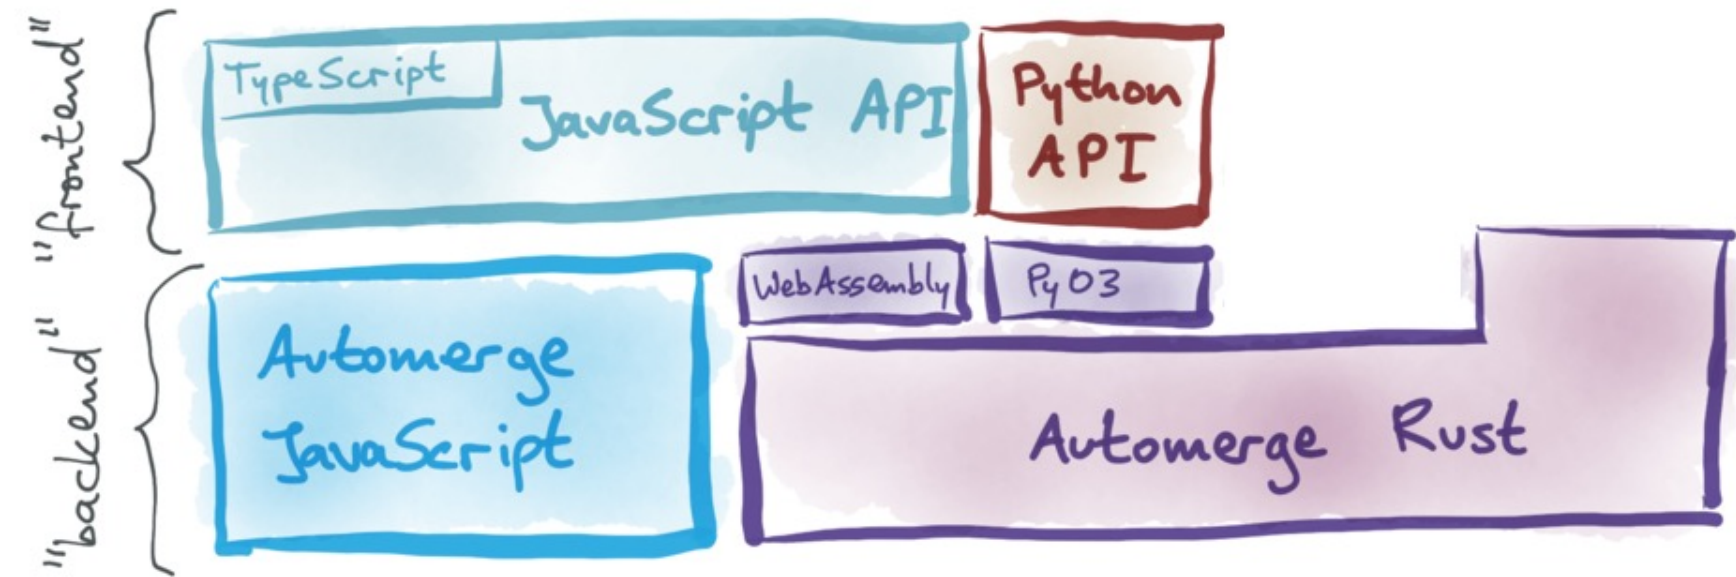
IMPLEMENTATION



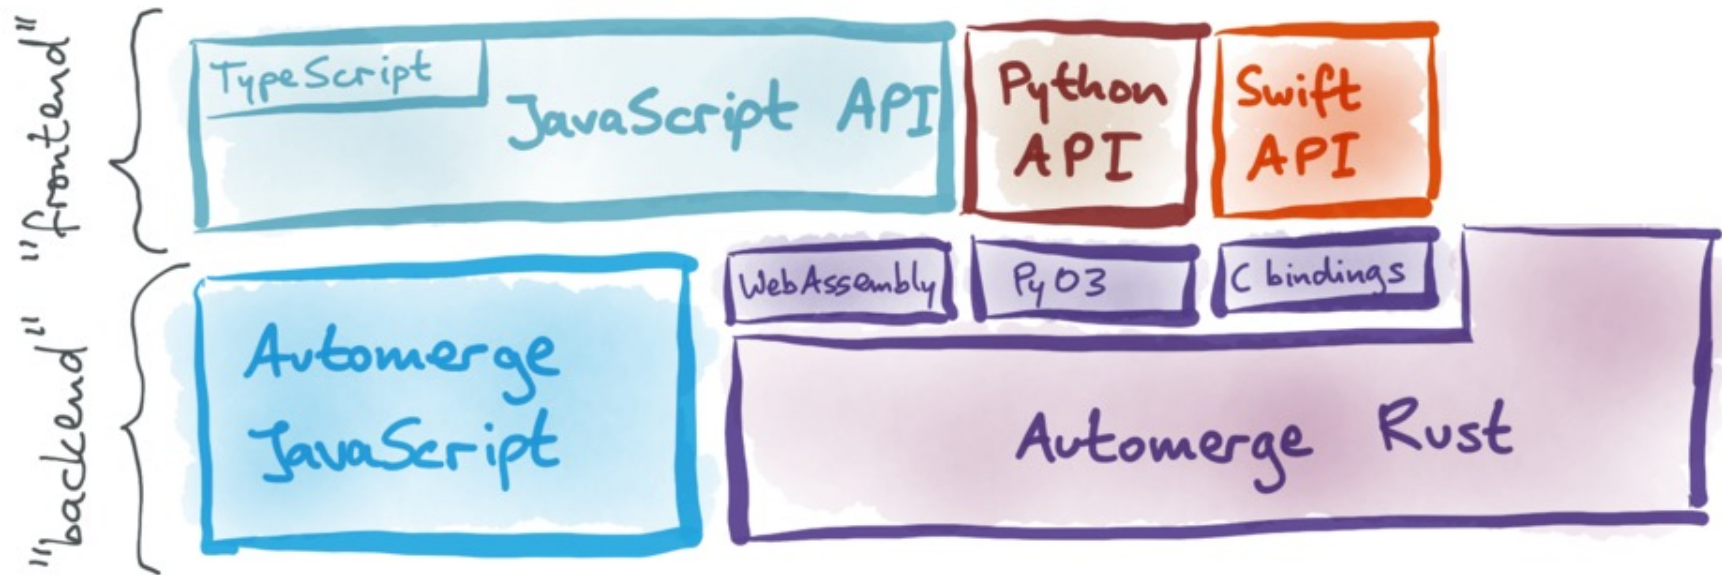
IMPLEMENTATION



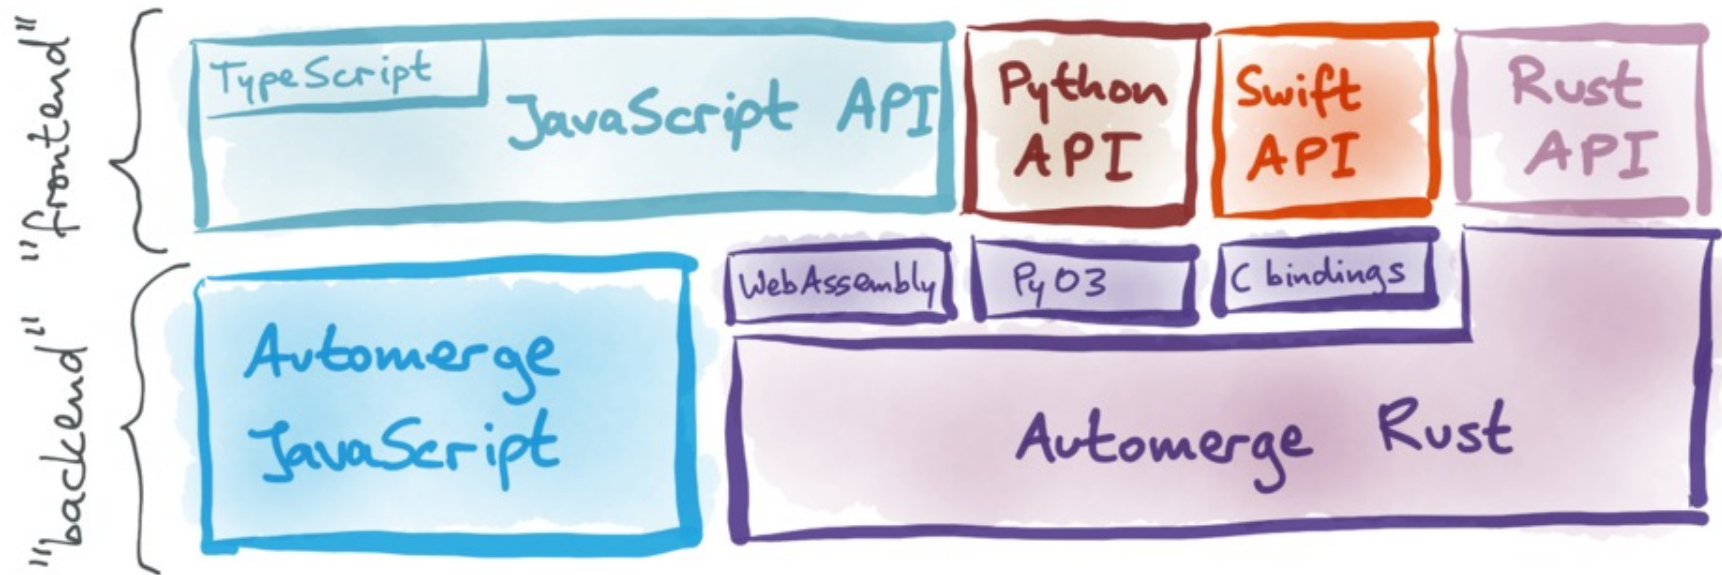
IMPLEMENTATION



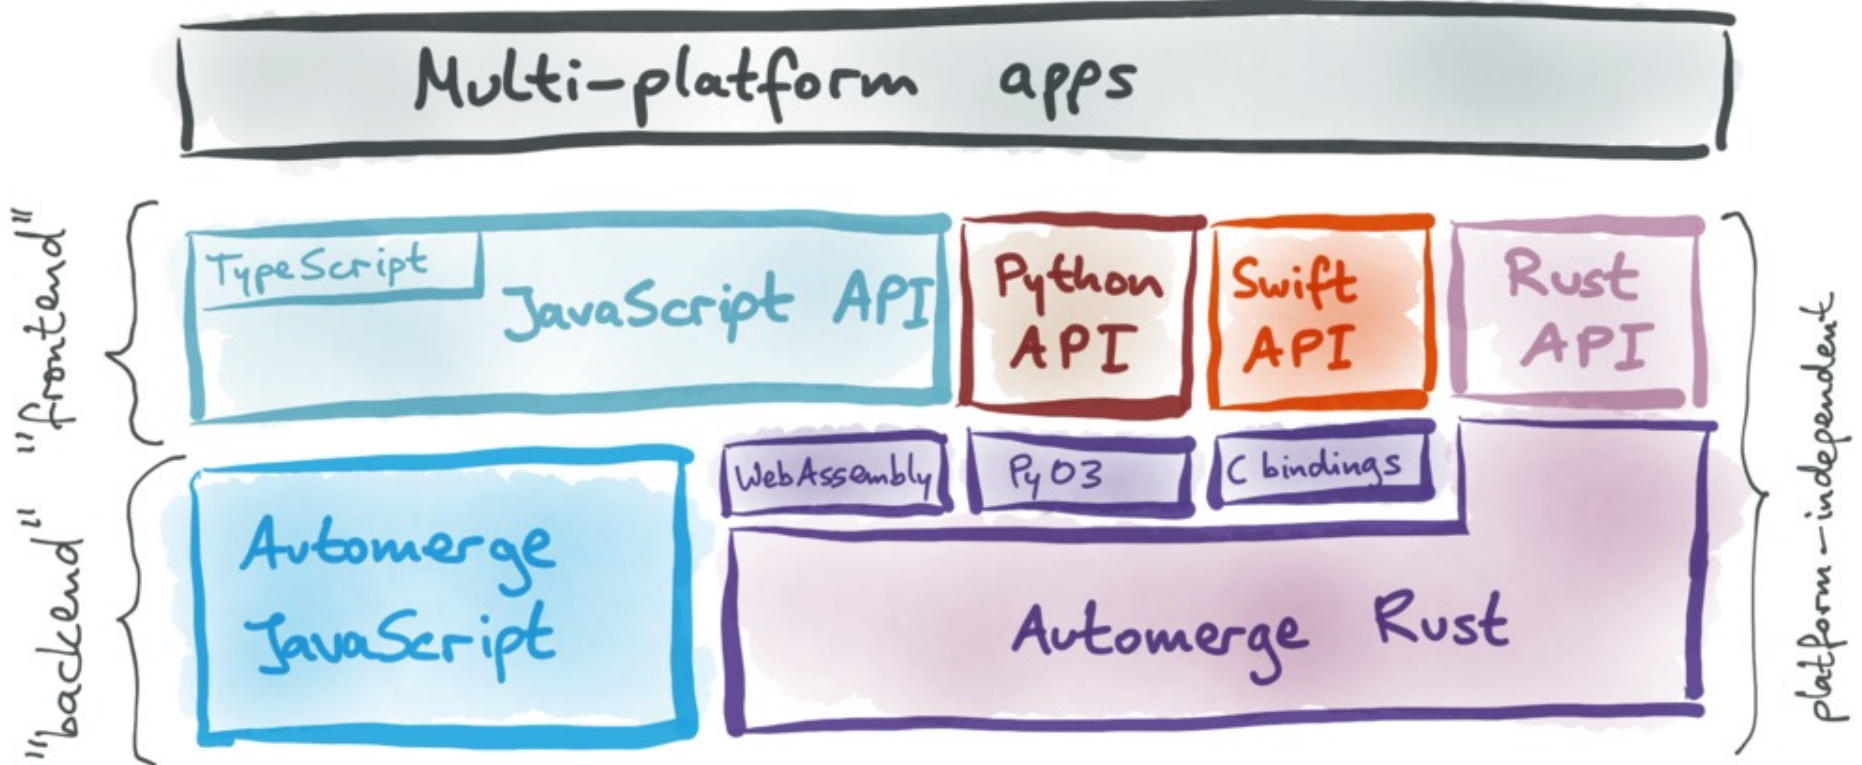
IMPLEMENTATION



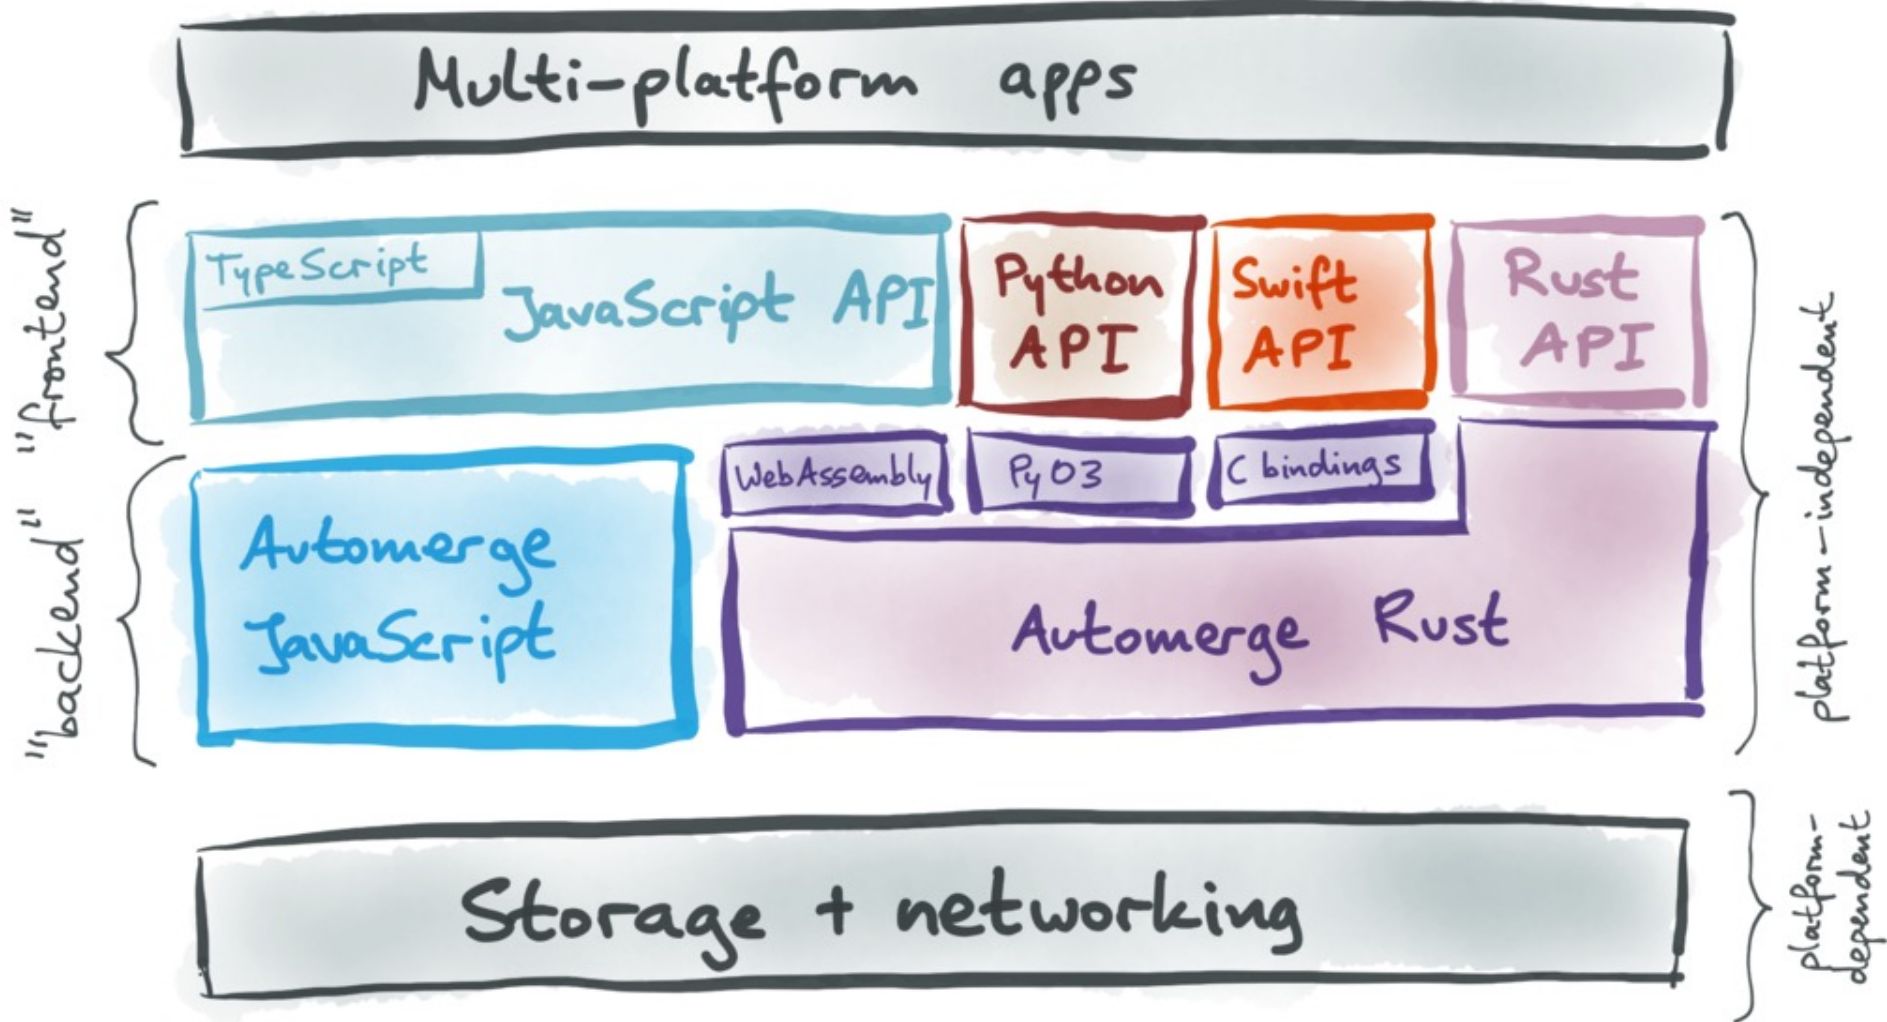
IMPLEMENTATION



IMPLEMENTATION



IMPLEMENTATION



Resources

Automerger	https://automerger.org/
Publications	https://martin.kleppmann.com/#publications
Email	martin@kleppmann.com
Twitter	@martinkl
Book	http://dataintensive.net/
Support me	https://www.patreon.com/martinkl

Huge thanks to my collaborators and the Automerge community and contributors, especially Peter van Hardenberg, Orion Henry, Alex Good, Andrew Jeffery, Rae McKelvey, Herb Caudill, & others!

LEVERHULME
TRUST



Ink & Switch