

ADVANCES IN REAL TIME

CHALLENGES AND SOLUTIONS IN INTEROPERABLE PAYMENT SYSTEMS

KOSTA PERIC, MILLER ABEL, AND MATT BOHAN

Finance can be very complex, but at its core it's about exchanging value between parties. In a word, payments. Making payments is a fundamental financial action for most people around the globe, whether they use cards, cash, or SMS to make them.

In terms of financial infrastructure, then, a payments platform is a fundamental piece, the anchor that tethers the daily financial actions of customers to the daily holdings of financial providers. Such platforms commonly have a national scale because of country-specific currencies and regulations. However, they also can be regional, as we see in the EU and the South African Development Community, and international, as we see with credit card systems. International platforms will become more necessary as digitization continues to put the common citizen in touch with others far beyond the borders of their home country.

In fact, the rise of digital financial services has changed the nature of payments and payment platforms in many ways. Customers can transfer money across great distances with a few taps on their mobile phones, with low or no associated fees. Providers can reach customers they weren't able to before—namely, people with very limited assets who are either

unbanked or underserved by predigital financial services. In emerging economies, telecoms and other non-bank institutions are stepping in to take advantage of this opportunity by offering digital payment services, such as mobile money.

This new dynamic between customers and providers, and the opportunities this evolution presents, has raised some compelling questions about how payment platforms exist and interact at the system level. Because of the positive impact digital accounts and payments can have on the lives of the poor, the Bill & Melinda Gates Foundation has explored these questions in great detail.

The outcome of that exploration was the development of the Level One Project (L1P), an initiative to promote the creation and evolution of inclusive, interoperable digital payment platforms. At the center of this initiative are the L1P principles (see box), which list the attributes essential to making payment systems accessible to poor customers and support-

ive of various financial service providers (see <http://leveloneproject.org> for more details).

In this article, we explore some of those principles and the challenges of implementing them. We also explore how modern technologies related to distributed ledgers were adapted to support a reference implementation (see below) of the L1P principles. Released in October 2017, this reference implementation demonstrates the viability of pro-poor payment platforms and offers governments and commercial providers a flexible way of developing such platforms.

Before we look at specific design principles and solutions, however, we should understand why they are important.

DIGITAL PAYMENTS AND FINANCIAL INCLUSION

Cash may be accepted everywhere, but it's often expensive to use, as informal couriers and lenders charge high fees. Nevertheless, people who don't have a lot of money still use cash constantly, and making transactions by hand is extremely risky and time intensive for them.

For someone who has only used cash, making payments in a formal manner using a banking or digital payment system—whether by check, card, or digital transaction—is nothing short of revolutionary. Time is saved. Funds are secured. And, in the best cases, fees are reduced.

High fees have long been one of the biggest obstacles to the world's poor making formal payments. With the advent of

ABOUT THE AUTHORS

Kosta Peric is Deputy Director for the Financial Services for the Poor initiative at the Bill & Melinda Gates Foundation, where in charge of payments. Previously he was the chief architect of SWIFTNet, SWIFT's global secure network that connects 10,000 financial institutions and corporations in the world, and co-founder of Innotribe, SWIFT's initiative to enable collaborative innovation in financial services. He is the author of *The Castle And The Sandbox*, a book on how to innovate in conservative companies using open innovation.

Matt Bohan serves as a Senior Program Officer at the Bill & Melinda Gates Foundation, with a specific focus on the Financial Services for the Poor. Matt supports the program's Level One Project, where he's working to bring mobile payment and other financial services to the two billion people in the world who live on less than two dollars a day. Prior to working at the Gates Foundation, Matt was Senior Director at Alvarez & Marsal.

Miller Abel is a Principal Technologist for the Financial Services for the Poor team at the Bill & Melinda Gates Foundation. He has more than 35 years of experience in the software and services sector, where he has held a variety of positions in both small and large systems companies. Miller specializes in fast-changing Internet and mobile businesses where ambiguity and uncertainty must be overcome by solid strategy and converted into business results. Miller has 30 patents issued and pending (not including foreign equivalents) in a variety of technology areas including contactless proximity communications.

© 2018 Kosta Peric, Miller Abel, Matt Bohan

innovations / volume 12, number 1/2

71

LEVEL ONE PROJECT PRINCIPLES

The attributes essential to making payment systems accessible to poor customers and supportive of financial service providers include:

- An open-loop system that is available to any licensed digital financial service provider (DFSP) in the country, including banks and licensed non-banks.
- Real-time and “push” payments, and payments that are irrevocable, which remove many of the risks and costs inherent in batch-processed and “pull” payment systems.
- A system that is governed by the DFSPs that use it and regulated by a government financial authority. This well-tested model creates a feeling of fairness among participants.
- A system that allows same-day settlement or better among participants.
- A system that operates on a “not-for-loss” or “cost-recovery-plus-investment” basis. This does not preclude DFSPs—or other service providers in the ecosystem—from earning profits through use of the platform.

See <https://leveloneproject.org>.

digital financial services, providers finally have a way to reach a large number of customers with very low overhead. For example, the cost of facilitating a digital payment from mobile phone to mobile phone is negligible, as are the fees providers need to charge to serve their customers profitably. Digital financial services finally give poor customers the speed and security of formal banking at a price they can afford.

This has helped drive a boom in mobile money products and providers across the world’s emerging economies. The first mobile money service hit the market in 2004, and by 2016 there were 277 active services and more than half a billion registered accounts worldwide.¹ In Uganda and other countries, mobile money customers outnumber traditional banking services customers.²

Building individual financial resiliency is a slow process, but mobile money has been around long enough to have a measurable effect on poverty. Last year, researchers from MIT showed that 194,000 households in Kenya emerged from extreme poverty because they had access to the M-Pesa mobile money serv-

ice.³ Other studies have shown that digital accounts help mothers buy healthier food for their families and help farmers save and invest in future harvests.^{4,5}

Today, mobile money is at an inflection point. Established providers continue to thrive, but new ones have a hard time entering the market. Person-to-person payments are common among the poor, who regularly lend and transfer money among friends and family, but other forms of payments are extremely important as well, such as government social disbursements and transactions with merchants. The integration of these kinds of payments into the mobile/digital milieu has so far been limited.

Cash still has one remaining advantage over mobile money: universality. Not everyone wishes to transact with mobile money, and parties wishing to transact may not use the same mobile money service. Everyone will, however, accept cash. This gives would-be customers a persuasive reason to hold on to their cash and think twice about signing up for digital accounts.

The current point of inflection, then, rests on interoperability among services.

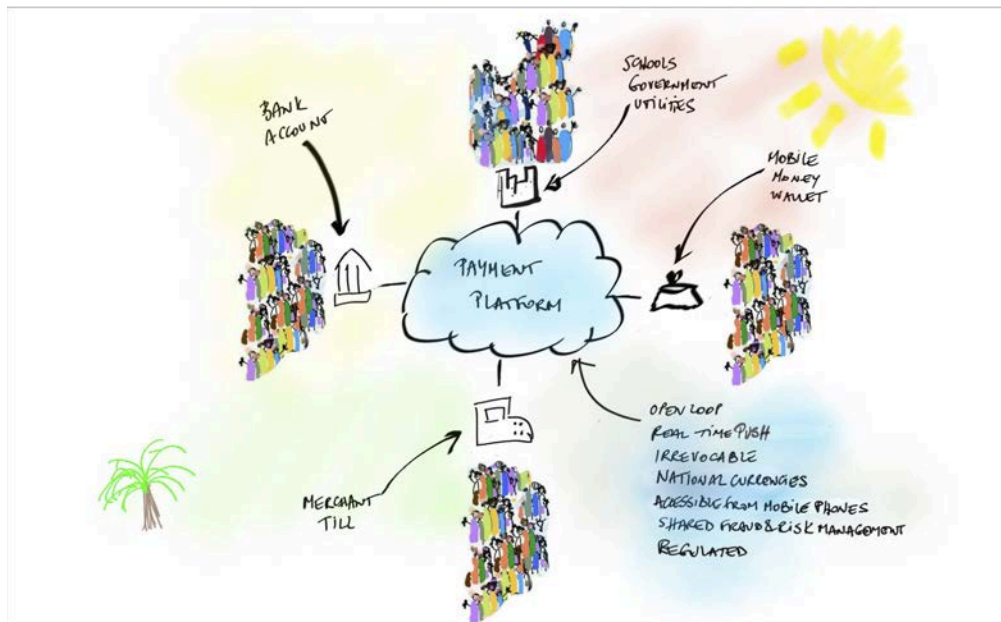


Figure 1. Level One Project Vision and Principles

For mobile money to continue its growth in terms of both the number of customers and the number and diversity of providers, services cannot continue to function as closed loops, wherein subscribers of one service are unable to transact with subscribers of a different one. For customers to get the full benefits of digital finance, they must be able to transact with anyone, not just the random cross-section of the population that happens to have an account with the same provider.

Interoperability will also help extend mobile money into merchant and bulk payments. With interoperability, merchants and employers will be able to use a single service to issue and accept money to and from all their customers and employees. Without interoperability, they must either choose one service, which excludes customers and employees using other services, or subscribe to them all. It's a similar dilemma for government-to-

person payments, though on a much larger scale, often millions of citizens each month.

Finally, interoperability based on a shared central service will make it easier for new providers to enter the market. Building a customer base is something every provider needs to do quickly in order to achieve solvency. It's difficult to do when they can only offer new customers the ability to transact with a few hundred or thousand fellow subscribers; most customers will prefer to sign up with established providers who already have sizable loops. If transactions were possible across all services, all customers would exist in the same loop, and new providers could compete based on cost and features rather than the size of their subscriber base.

The L1P principles offer a vision for how an interoperable payments platform can serve the poor. Figure 1 shows this

vision, along with the underlying design principles.

BUILDING L1P-ALIGNED PLATFORMS: KEY BARRIERS TO OVERCOME

Building an interoperable payment platform on a national or regional scale is no small feat. It requires focusing clearly on the ultimate objectives, providing strong governance and oversight, and having a solid understanding of the barriers ahead.

From our experience so far, interoperability poses two key technological barriers that projects need to overcome: the lack of efficient clearing and settlement of payments across different financial service providers, and the lack of standardization in connecting these providers to a common platform. The L1P principles and reference implementation address both.

In terms of clearing and settlement, L1P offers three basic principles: (1) payments should be initiated and authorized by the payer not the payee; (2) payments should be completed and verified in real time; and (3) payments should be irrevocable. The first principle specifies that as with cash, an end user initiates a payment by pushing value to the receiver; and, unlike direct debits or checks, money never leaves an end user's account without direct authorization at the time of the transaction. The second principle states that when one person sends money to another, the sending account is debited and the receiving account is credited simultaneously and immediately, whether or not the counterparties use the same provider. As with the first principle, this is necessary for making digital money function similarly to cash, which will make it a meaningful and comfortable user alternative. The third principle states that payments are final when they are cleared; they can't be recalled by the person mak-

ing the payment. This is a hallmark of cash that is essential to adoption.

By contrast, the most common example of a pull payment is a check: when processing a check, the receiver's account requests that money be pulled from the account of the sender, who has pre-approved this request by signing the check. Checks are not cleared immediately. It often takes several days from the time the commitment to pay is made and when funds are credited to the receiver. And during this time before the check is cleared, it remains revocable because the sender may stop payment of the check before it clears. In this case, the request for payment is revoked by the payer—in other words, the check bounces. Moreover, because a pull payment like a check is a request for money and is not cleared in real time, there may not be adequate funds in the payer's account to fulfill the request. The sender is likely charged a fee for pre-approving a payment that couldn't be fulfilled.

Push payments are irrevocable because the sender's provider will only push money that is available in the sender's account. There is no way to cause an overdraft. And because push payments are cleared in real time, there is no ability to revoke. In short, the payment becomes final when it is authorized.

L1P principles prescribe these payments because they are technically simpler, which makes them more affordable, and because they are immediate. With pull payments, the receiver has to wait while the sending account verifies that the request for money is legitimate and can be fulfilled. With push payments, sender and receiver get instant confirmation. Making payments irrevocable enables them to happen in real time. However, to achieve real-time payments in an interoperable system, another issue must be addressed: the timing of transactions between providers.

When a user of one service sends money to a user of another service, one service provider ends up in debt to the other. Both customers' accounts must be credited and debited, as must the ledgers belonging to each provider. Therefore, a payment transaction must occur at some point between the providers.

Settlements between financial service providers already occur within the established banking systems in every country. Banks in this system usually settle their transactions with each other once every business day. This means that all the customer transactions for the settlement period are tallied up and cleared at once, in a single bulk transaction. This is another case where customers must sometimes wait for funds to be available, which results in payments that are not real time.

This can be a problem for mobile money and other non-bank providers in an interoperable system because they also depend on these once-a-day deposits with banks. To guarantee all the digital money reflected in the accounts of their end users, non-bank providers are required to have an equal supply of "real" funds saved with a bank so that, for example, if a provider's customers all decide to withdraw their funds on the same day, that provider is able to give each one their full account value in cash.

In an interoperable system, non-bank providers need an additional reserve of funds specifically designated to cover cross-platform transactions. When a customer sends money to someone who uses a different provider and the first provider becomes indebted to the second provider, that debt is paid using this designated account.

To keep pace with its users' cross-platform transactions, a provider will periodically move money into its designated account. While these deposits typically settle only once per business day, customer transactions happen much

more frequently. This means that the amount of digital money transacted across platforms by customers and the amount of cash a provider has on deposit to cover these transactions can easily fall out of sync. In other words, if there is a sudden surge in cross-platform transactions over the course of a day, one provider could end up owing another provider more than it has set aside to pay.

One solution is to make a customer wait for cross-platform transactions to clear. As stated earlier, though, this is unacceptable, because it fails to mimic the performance of cash and thus fails to meet the customer's needs.

Therefore, the time window for settling provider-to-provider transactions must be reduced. In systems aligned with L1P principles, this time window is reduced from a full business day to hours or even seconds. Provider ledgers are therefore updated and kept accurate multiple times throughout the day, thus reducing the capital requirements for supporting interoperability.

The other key barrier we encountered in deploying the L1P principles is the lack of appropriate standards to connect a financial service provider's system to the payment platform. Systems aligned with these principles fall into the category of real-time retail payments (RTRP) systems. Other systems—such as real-time gross settlement systems, which cover bank-to-bank payments that are high in value but low in volume, and automated clearing houses, which deal with low value-batch payments—traditionally come from the banking world and tend to be served by messaging standards, like the ISO 20022. The RTRP space is more technologically demanding, and platforms aligned with L1P principles tend to connect a variety of providers in addition to banks, such as mobile money providers, microfinance institutions, payment aggregators, and merchant networks.

Because of the technological demands and diversity of players, an application programming interface (API) design principle offers a viable solution. APIs can accommodate the deeper and more efficient integration required by a high volume of transactions while placing a relatively low development burden on individual providers—certainly much lower than constructing their own platforms or connections.

However, there is as yet no pervasive standard for this type of integration. Therefore, one challenge to overcome in deploying platforms aligned with L1P principles will be to settle on an API standard that will enable easy and rapid integration of financial providers.

MOJALOOP: THE OPEN-SOURCE REFERENCE IMPLEMENTATION OF L1P

As efforts to deploy systems modeled on L1P principles in actual country markets progressed from 2013 to 2015, it became apparent that the barriers described above were slowing them down. Deployments require expertise in designing and constructing payment systems—a resource that is difficult to procure, especially in the African and Asian countries where L1P principles were being considered.

We realized that we needed to support these deployment efforts by providing additional guidance and examples of plans and designs, and even of software code, to make sure the projects could progress according to plan. The team then commissioned a reference implementation for L1P, which took the form of open-source software that embodied the architecture and design principles of L1P and demonstrated how to overcome the barriers described earlier.

Publishing the reference implementation as an open-source asset was important for a few reasons. Since its inception,

all materials and knowledge gained from the L1P principles have been shared on the project's website.⁶ This is consistent with the open, collaborative spirit that led to the development of these materials, and with the goal of enabling anyone to use the assets and to create L1P-aligned platforms without direct involvement with the Gates Foundation. As open-source software, the reference implementation remains available for adaptation and adoption, like all previous L1P assets.

Called Mojaloop—building on the word “moja,” which means “one” in Swahili—the software was developed by Ripple, Dwolla, The Software Group, ModusBox, and Crosslake Technologies. In October 2017, it was made available on GitHub and on its own website (mojaloop.io) under the Apache 2.0 license.

Mojaloop covers many of the essential uses end users find for real-time payments: person-to-person, merchant point-of-sale, payroll and bulk payments, multiple accounts and users, and fraud monitoring. Four central components enable these uses:

- A central directory service, which routes each payment to the correct service/provider in the ecosystem
- A central ledger service, which tracks transactions for compliance and settlement among providers
- A fraud service, which allows providers to share transaction information to validate and secure payments
- A central rules service, which sets policy across the system

Of special interest in this article is how Mojaloop enables faster clearing and settlements between providers. There were several iterations in the architecture and design. The first of these considered a public blockchain. The idea was that the clearing and settlement would occur as transactions on the single distributed

ledger connecting all financial service providers in the ecosystem. The design accommodated the actual movement of money between providers (i.e., direct settlement using central bank money), which is clearly superior to managing settlements outside the platform. The design also could assign the job of clearing to the distributed ledger, with settlement still occurring on a periodic basis (using a settlement bank) for the regulatory environments where this is mandated.

While it satisfied our objective of faster settlement, this architecture presented two key drawbacks when considered in the context of a national payment platform:

- **Lack of national-level data control and confidentiality.** Because it used a public blockchain, the design would replicate national payment data across the entire Internet. This obviously conflicts with the natural sovereignty regulators typically require. A country-level blockchain could have been considered as an alternative, but this was not pursued, as other more practical solutions exist (see below).
- **Lower capacity for transaction volume.** The design could accommodate hundreds of transactions per second (depending on the nature of the blockchain used), which is impressive, but it falls well short of the thousands per second required for real-time retail payment platforms.

The team thus looked for other solutions. In the end they chose to use the Interledger protocol, along with specific code implementing a central ledger in the platform.

The Interledger protocol is itself an open-source specification, invented by Ripple and developed through a broad collaboration under the World Wide Web Consortium.⁷ The protocol enables cryptographically assured ledger synchroniza-

tion. The major advantages of the protocol are the low level of requirements to conform with it, and the extreme efficiency and scale in processing transactions. These aspects make it especially useful in high-volume, low-value retail payments systems—that is, in systems that millions of poor customers will use to make lots and lots of transactions in very small amounts.

Mojaloop's current design places a central ledger within the payment platform and uses the Interledger protocol to synchronize the ledgers of all the financial providers involved in a given transaction. This securely synchronizes payment transactions within the system. The design provides for immediate settlement between providers or a separate clearing and settlement in a separate bank at very high transaction rates—meanwhile preserving the privacy and sovereignty requirements of the entire system.

The second barrier mentioned above was the lack of standards for connecting financial service provider systems with the payment platform. The Interledger protocol provides easy use and fertile ground for such a standard. Thus, while the Mojaloop platform was being designed, the team at the Gates Foundation engaged another team of developers to work on an API.

The mobile money platforms of Huawei, Ericsson, Mahindra Comviva, and Telepin provide the majority of mobile money operators with the key technologies needed to meet their customers' needs. These four companies agreed to collaborate on creating an API standard for industry interoperability. They also agreed to upgrade their products to meet this standard, which automatically gave all their customers the option of using it. The latest update to the API (available at mojaloop.io.) provides a layer of semantic interoperability and can be deployed over the Interledger protocol.

As stated earlier, everything produced or enabled by L1P is available in the public commons. While the essential challenges are the same, every market will have its own regulatory and technical environment. The assets aligned with L1P principles are designed to be flexible, and to reflect the needs and intentions of their adopters rather than their creators.

CONCLUSION

Numerous governmental and nongovernmental bodies have codified principles of financial inclusion. These include the Maya Declaration from the Alliance for Financial Inclusion and the High Level Principles for Digital Financial Inclusion put forth by the G20. The Level One Project principles are unique in that they are specific. Rather than pointing to the broad strokes necessary to make progress, such as infrastructure and government leadership, the L1P principles identify the particular aspects and functions a digital financial platform must include in order to be inclusive.

Technology inspired by distributed ledger technology can help resolve the challenge of real-time settlement across an array of diverse, interoperable providers and ledgers. This is borne out in the Mojaloop example, which applies the Interledger protocol to achieve multilateral net settlement in a digital financial ecosystem consisting of mobile money providers, banks, and merchant account holders.

By laying out key components, the L1P principles provide a valuable tool for increasing and sustaining digital financial inclusion. By manifesting these components, with the help of DLT-inspired strategies, Mojaloop and the mobile money API make the vision of an interoperable ecosystem much more possible to realize.

There is potential for such ecosystems to become common around the world, starting in the emerging economies, where the size of the unbanked population and the zeal for financial innovation are both quite high. From there these ecosystems can blossom into regional and even continental platforms. As they do, Mojaloop may again be useful in solving the new challenges that arise, such as how to facilitate payments across borders and currencies.

¹. GSM Association. 2017. "State of the Industry Report on Mobile Money: Decade Edition (2006-2016)." Available at https://www.gsma.com/mobilefordevelopment/wp-content/uploads/2017/03/GSMA_State-of-the-Industry-Report-on-Mobile-Money_2016-1.pdf.

². World Bank Group. 2014. "Global Index." Available at: <https://globalindex.worldbank.org/>

³. Suri, T, Jack, W. "The Long-Run Poverty and Gender Impacts of Mobile Money." 2016. *Science*, 354: 1288-1292. Available at <http://science.sciencemag.org/content/354/6317/1288.full>.

⁴. Doepke, Matthias, and Michèle Tertilt. "Does Female Empowerment Promote Economic Development?" Centre for Economic Policy Research discussion paper no. 8441, June 2011.

⁵. Brune, Lasse et al. Facilitating Savings for Agriculture: Field Experimental Evidence from Malawi. NBER working paper number 20946, February 2015.

⁶. See <http://leveloneproject.org>.

⁷. See <http://Interledger.org>.