Abstract

The method of billing users for mobile telephony is based on systems developed over time for fixed networks. We survey the technology involved, and argue that these systems will become increasingly inadequate for large populations of mobile users where frequent roaming is involved. We present two micropayment schemes which permit a caller to inject a payment stream into the network which allows multiple network operators and value-added service providers to be paid in real time. The methods support dynamic pricing by the association of a pricing contract with the call which specifies the cost for each leg of the call route. The system will alleviate problems of mobile fraud, eliminate the need for interoperator billing agreements, and simplify payment for value-added network services. We discuss the relative merits of the two systems described and the characteristics of the prototype implementation.

Flexible Real-Time Payment Methods for Mobile Communications

MICHAEL PEIRCE AND DONAL O'MAHONY, TRINITY COLLEGE DUBLIN

In the fixed network, customers have an enduring relationship with their network operator. Payment takes place after services are used, and the fact that a particular local loop is involved means that additional customer authentication is unnecessary. Once the user becomes mobile, many additional problems come into play. Operators of current mobile networks have had to face this problem, and have responded using billing schemes based largely on extensions to existing fixed network billing systems.

Cellular network fraud has become a major problem, and efforts to curb the extent of this fraud have led to the imposition of service restrictions on users, particularly in regard to roaming. In the future, when the number of mobile users of either cellular networks or personal mobility services exceed the number of fixed users, the existing approach to billing will not be sufficient.

In this article we outline a real-time payment scheme that allows users of mobile communications services to arrive in a new network, and avail themselves of network service using real-time payment to both the primary network operator and any other service providers that may be involved in the call.

Existing Approaches to Billing

The practice of recording the details of individual calls for billing purposes has been in use since the commercial deployment of early manual telephone exchanges in the late 19th century. In those times *call detail record* (CDR) information for long distance calls was collected and recorded by cordboard operators at the exchange [1]. The operator manually wrote the details onto a specially formatted record called a *toll ticket*. These tickets were later sent to a clearing office where customer bills were generated.

The same basic principles for charging for the use of telecommunications networks are in place today. The telephone exchanges have evolved into complex digital switching systems totally controlled by software. When a subscriber places a call the local exchange automatically creates a record of the call details [2], a process known as *automatic message accounting* (AMA) or *toll ticketing* (TT). The CDRs are stored in a file at the local exchange and periodically sent to a centralized billing system, usually at another location. This offline procedure can vary from physically transporting magnetic tapes to transmitting the records across a data network. If the records are transferred immediately, the process is known as *hot billing* or *online charging*. This allows bills to be processed on request or within a given time limit. However, it is not yet in widespread use among operators.

Billing software extracts information from the CDRs and calculates the cost of the call for the customer by applying price rate tables based on called distance, call duration, time of day the call was placed, and subscriber type (residential or business). Every billing period a bill, possibly detailing the calls placed, is sent to the customer for payment. The billing lifecycle is summarized in Fig. 1.

The contents of a CDR vary from operator to operator. This is partly due to their use not only for billing, but also for the measurement of traffic and quality of service. Switch architectures from different manufacturers also produce information in different formats. Recent standards for call records do exist [3–7], but are not yet widely adhered to. However, there are some critical fields for billing, usually present in some form. These are listed below.

Source of Call — The CDR must uniquely identify the party to charge. Normally calls are billed to the owner of the phone from which a call is placed, identified by the calling line number. Other scenarios might bill the called party number (reverse charging), an account number (calling card or credit card services), or a personal user identity (allowing personal mobility).

Destination of Call — Usually the called number (digits dialed) and a translated number if appropriate. Translated numbers arise from call forwarding and special rate numbers (toll-free, premium rate, local national rate). The distance rate to apply for billing is obtained from this information.

Time and Duration — These are the date and time the call took place and the length of time it lasted.

Routing Points — The identity of resources and equipment, such as trunk lines or gateways, used during the call. Distance billing may be partly based on the routes taken by the call, due to the interconnection agreements between different net-

1070-9916/99/$10.00 © 1999 IEEE

work operators for resource usage. Traffic analysis is also based on routing details.

Service Information — The type of call service used, such as basic call, premium rate, toll free, conference call, or other supplementary services. It may also include the quality of service (QoS) provided, such as use of a low-delay optic fiber link instead of a slower satellite connection.

Charging and pricing information is usually applied when the records reach the billing system. The size of a CDR can range from 20 bytes to several hundred bytes, but for performance efficiency it should be kept to a minimum. If a call is not answered, a CDR may still be created, although the customer will not be billed for these. Sometimes two or more records are created for the same call, an originating record and a terminating or trunk record. This is useful when all the necessary information is not available at one place in the network.

The only proof that a call took place is the CDR, and sophisticated duplication schemes are used to ensure that this data is not lost. The raw data generated by the switch is also kept for the purpose of settling disputes. However, there is no mechanism in place to prove the authenticity of the data. The call records can be denied by the customer or falsified by the operator.

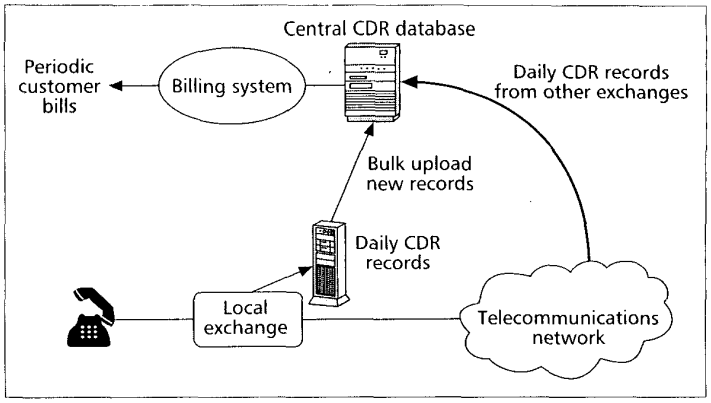

■ **Figure 1.** *The billing cycle in existing telecommunications networks.*

Billing in Mobile Networks

Rapid growth in mobile communications has given rise to a large number of independent network operators, spanning many different geographic areas and countries. When these operators use a common mobile standard, it is possible to allow subscribers to roam from the home network to a visited location, choosing between the new operators available. The Global System for Mobile Communications (GSM) [8], with over 350 network operators worldwide, provides more than 215 million customers with the ability to make and receive calls when outside their home network. In order to coordinate the initial deployment of the GSM standard, 15 mobile network operators signed a memorandum of understanding (MoU) committing to introduce GSM systems by 1991. This early agreement has evolved into an international association of GSM network operators, the GSM Association [9]. Its purpose is to guide the commercial development of GSM, while working alongside the technical standards bodies such as the European Telecommunications Standards Institute (ETSI).

Billing and accounting procedures arising from international roaming are regulated by the GSM Association. It states that charges made to a roaming subscriber in a visited network must be collected by his home operator with whom he has a subscriber agreement [10]. When the visited network operator provides services to the subscriber, they need to be assured that the subscriber's home network will compensate them for the charges incurred. A mobile user has no contract or relationship with the visited operator. For this reason roaming is only permitted between networks which have arranged a *bilateral roaming agreement*. The GSM Association Billing Administration and Roaming Group (BARG) has laid down the regulations applicable to such roaming agreements.

The visited network operator provides services to roamers without the need for additional subscriber contracts or credit authorizations. To ensure that the subscriber is genuine, his/her home location register (HLR) is contacted. Authentication is performed based on the knowledge of a shared secret between the HLR and subscriber. Upon a valid identification the subscriber can make and receive as many calls as desired through the visited mobile network. CDRs are generated by the visited network and later forwarded in bulk to the home operator for settlement. The network operators settle payment for the resources used between themselves, and the charge is ultimately reflected in the user's bill.

The Transferred Account Procedure (TAP) [11] details how the CDRs of roamers should be transferred from the visited network back to the home network. The file formats and transfer methods between operators are specified by the GSM Transfer Account Data Interchange Group (TADIG). Electronic Data Interchange (EDI) is used to exchange the TAP files in a standard message format. This is often performed using file transfer over an X.25 network.

Mobile networks not based on GSM technology use procedures other than the GSM TAP to transfer the call records of roaming subscribers. The Cellular Intercarrier Billing Exchange Record (CIBER) [12] is used for roamer billing throughout North and South America. CDRs of roamers are converted into the CIBER format before being exchanged. The CIBER standard is published by Cibernet, a provider of financial settlement services for wireless operators.

Multiple CDRs, or toll tickets, can be generated for different legs of a GSM call. These include the originating, terminating, and roaming call components. In a mobile-originated call the calling party usually pays for all stages. When a roaming mobile is called, the subscriber pays for the component of the call from the HLR to his location in the visited network. With optimal routing [13], where calls are routed directly instead of via the home network, the roaming component can be minimized. Figure 2 shows mobile subscriber A, in his home network, calling a roaming subscriber B. A will be billed for the originating leg (TT1), while B will be billed for the roaming leg (TT2) and terminating leg (TT3).

There are 52 different fields within a GSM toll ticket [3]. The International Mobile Subscriber Identity (IMSI) identifies the source of a mobile-originated call. In addition to regular CDR fields the toll ticket contains location area, cell ID, international mobile equipment identity (IMEI), and radio channel allocation.

The roaming agreement is used to establish trust between independent operators. However, there is no guarantee of incontestable charging or payment for any of the parties involved. In addition, each network operator may have to exchange CDRs and payment with up to several hundred other network operators. Such direct reconciliation is both costly and inefficient. In order to minimize the number of transactions, and hence the cost of settlement, a central clearinghouse or broker can be used. In the absence of hot billing, fraudulent calls may last for hours or even days before they are detected when the toll ticket is finally cleared. To reduce such fraud some operators will terminate all calls that exceed a specific time limit.

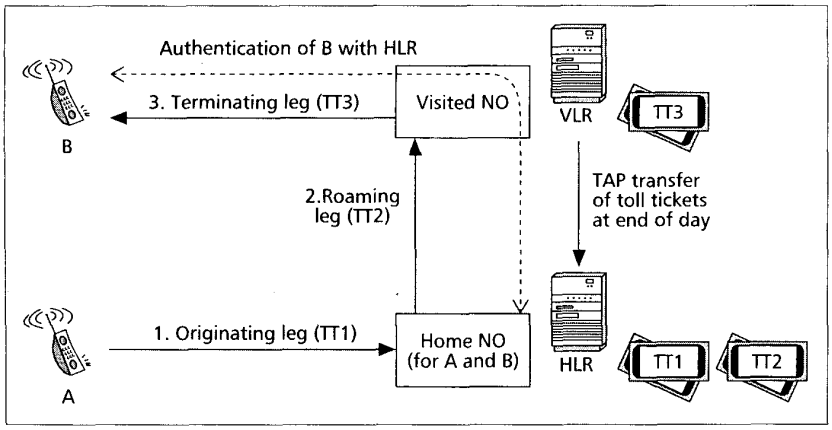

■ **Figure 2.** *Generation and exchange of billing records in the GSM system.*

Other problems such as privacy issues also exist. Use of CDRs results in databases being kept with details of the location and duration of every call a user makes or receives.

Retail and Wholesale Rates

The fixed network is based on a dual price system. For each call the user is charged a retail price, called the *collection rate*, by the originating network operator. In turn, the originating network operator is charged a reduced wholesale price by the terminating operator for completing the call connection. For a domestic interconnection between two operators within the same country, this wholesale price is known as a *termination charge* and is based on the cost of delivering the call to the final destination within the termination network. Thus, the charge may differ depending on what part of the network the call is going to or what resources are used.

The current situation with international interconnection works differently. The originating and terminating international operators agree on a wholesale price, called the *accounting rate* [14], for delivering traffic over their part of the international link. Traffic usage of the link is recorded; if there is an imbalance in the volume of incoming and outgoing traffic, the originating operator which generates more traffic pays the difference, called the *net settlement payment*, to compensate the terminating operator. The *settlement rate* is usually half of the accounting rate, which assumes that the cost of terminating the call is the same for each partner.

The accounting rate system has come under criticism and is currently being reformed. It was originally designed for an industry structure based on national monopoly providers and is biased against countries which send more calls than they receive. International calls to mobile networks introduce further problems. The settlement rate is set on the basis of fixed network termination costs. If the international call must be routed into a mobile network, with an additional domestic termination charge, the settlement rate may not be high enough to cover this charge as well as the costs incurred by use of the international facilities. Hence, one or more of the operators involved could conceivably lose money.

For both domestic and international interconnection the seconds of traffic terminated for each call are recorded in CDRs. These are added up at the end of the billing period and charged at the wholesale price to the originating operator. The GSM Association have introduced a wholesale tariff between GSM operators for roaming services, called the Inter-Operator Tariff (IOT). Previously roamers were charged the retail tariff of the visited network, which could change significantly with currency fluctuations. GSM operators continue to bill each other based on CDRs exchanged using TAP.

Prepaid Solutions

To increase the probability of receiving payment, the above billing schemes all rely on a strong legally binding contract established with the user. Where such a contract is not desirable, prepaid solutions allow the user to access specific services, for which they pay in advance. Calls are cut off in near real time when the prepaid amount has been used up, thus preventing a possibly unexpected large bill from accumulating.

The coin-operated payphone was one of the first prepaid solutions for the telephone network. Problems due to theft of the deposited money, the cost of collecting the coins, and counterfeit fraud led to the development of card-operated payphones. Prepaid cards, usually based on memory cards or smart card technology, are purchased from a distributor. As the card is used its value is decremented locally, often in response to toll pulses sent from the exchange. There is no need to verify the card online with a central database, as is the case when credit cards are used in payment.

Calling cards offer a temporary account with a network operator against which calls can be made. Such accounts can be prepaid or credit-based. In the former case the account status needs to be monitored and the call terminated in real time when the value is used up. International discount calling services are similarly account-based. They allow the use of an alternative network, and its reduced tariffs, by connecting to a user and presenting them with call-out services from that alternative fixed network.

Prepayment becomes more complicated in mobile networks because it is still desirable to be able to receive calls and roam internationally using many different network operators. The majority of prepaid mobile solutions are based on temporary accounts maintained at the HLR. Ericsson [15] is one supplier of such solutions in Europe. Its range of GSM prepaid solutions are intelligent-network-based with hot billing to allow automatic call termination when the account value reaches zero.

While many network operators now offer such services, only a handful allow prepaid international roaming. The solutions are *ad hoc* and are often priced to cover the highest possible tariffs in a visited network rather than the actual one in place. Some systems require a credit card against which to bill roaming charges, another uses a flat-rate call back service through the home network, and another requires that the GSM short message service (SMS) be used to send the desired number to the home network for call completion. Each solution usually has a substantial initial charge to cover costs of registering and maintaining identities in the HLR, making it infeasible for use by a casual visitor. There is clearly a need for improved solutions and agreement between operators to allow prepaid roaming in mobile networks.

Proposed Future Systems

In future mobile systems it is envisaged that there will be a large number of separately administered access networks, each with connections to one or more fixed networks. Through these users will be able to access a large variety of online services provided by an even larger number of competing *value-added service providers* (VASPs). Within any one region there might be hundreds of independent network operators, giving access to information and services provided by both local and

remote VASPs. The current billing and payment mechanisms outlined earlier, with implicit trust relationships between parties, will no longer be suitable in this scenario.

In Europe the name given to the future third-generation system is the Universal Mobile Telecommunications System (UMTS) [16]. The Advanced Security for Personal Communications Technologies (ASPECT) project [17] was a three-year Advanced Communications Technologies and Services (ACTS) project which investigated secure billing for UMTS applications, among other security features. ASPECT demonstrated an incontestable charging procedure, designed to allow small payments for value-added services [18, 19].

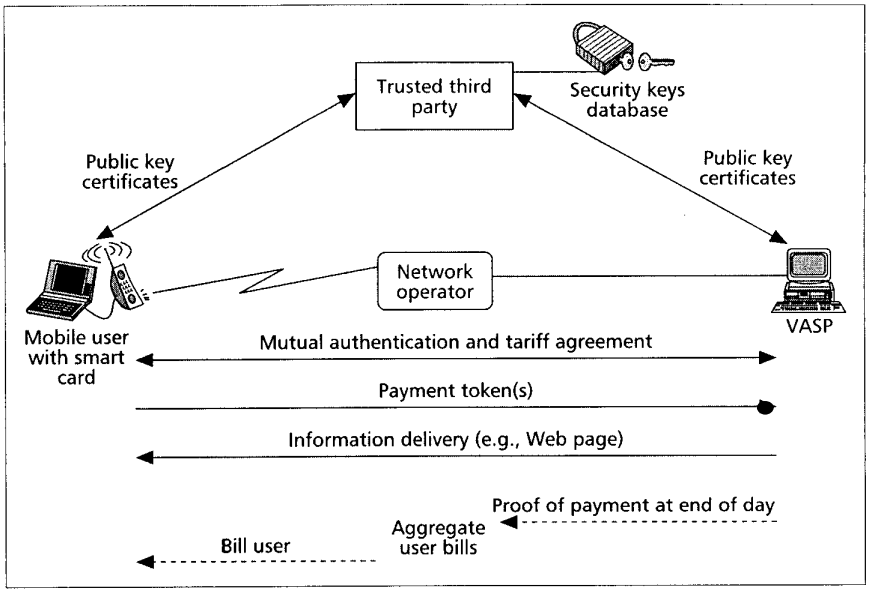

■ **Figure 3.** *Secure billing in ASPECT.*

The ASPECT approach was to break a call into two chargeable components. The first component was the basic charge for bearer services, the transport of call data, provided by the network operator. This is handled using traditional billing. The second chargeable component was the premium rate charge for use of services provided by a VASP. The ASPECT solution allows the mobile user to make many small payments directly to the VASP as the services are provided. The scheme is outlined in Fig. 3. Each payment token can only be generated by the user and is proof that he agrees to pay the VASP a small fixed amount. At the end of the day the VASP forwards the payment proof to the user's UMTS service provider, who then bills the user in the traditional fashion.

ASPECT improves on current solutions by providing incontestable charging for premium rate services, guaranteeing that the bill from a VASP is genuine. However, it still does not address network operator billing; nor does it guarantee payment from the user. The ASPECT project was completed in 1998, and a new ACTS project entitled UMTS Security Architecture (USECA) started. It is defining a complete UMTS security framework for standardization. In December 1998 an international consortium of telecommunications standards bodies, known as the Third Generation Partnership Project (3GPP), was formed to produce technical specifications for a third-generation mobile system. Both 3GPP and the GSM Association are also considering the requirements for charging and billing, based on CDRs, in third-generation systems [20, 21].

Micropayment Technology

In order to prevent fraud there is a need to be able to pay in real time for telecommunications services. Traditional billing allows the total amount to be paid afterward, but cannot ensure payment or provide incontestable charging. The entire amount cannot be paid in advance since the duration of the call or quantity of services used is not usually known beforehand. Many of the prepaid solutions discussed earlier work on the principle of making several small payments throughout a call. For example, this might involve depositing coins into a payphone or deducting units from a prepaid card at regular intervals. However, these methods only allow a single specific network operator to be paid. With mobile communications a roaming user is likely to use many different network operators and VASPs. It is desirable to be able to make efficient repeated payments of small amounts, called *micropayments*, to all the parties involved in a call as the services are used.

Many methods for electronically paying for services or goods across a network have been proposed [22]. They include the secure use of credit cards, electronic checks, digital cash, and subscription-based services. Each of these *macropayment* instruments have a minimum transaction overhead, usually imposed by the issuing bank, which prevents them being used for payments of a few cents. This is especially true of credit cards and electronic checks. A second prohibiting factor is their heavy use of computationally expensive cryptographic operations, such as public key cryptography. Digital cash systems, which try to mimic some of the properties of real cash, usually employ some form of digital signature technology for every transaction. For example each electronic coin might be signed by the issuing bank, or signatures might be used for mutual authentication between parties. These factors make macropayments too inefficient for frequently repeated transactions such as making a payment per second for a telephone call. In contrast, *micropayment* solutions are designed to allow efficient transfer of very small amounts, perhaps less than a penny, in a single transaction.

Micropayments allow new opportunities for charging, not only for basic transport services but especially for premium rate services and value-added information. Mobile users might pay to receive weather information, view financial market data, join a conference call, or listen to their voicemail. To be viable such a scheme must be able to perform a large number of transactions per second at a very small cost. This is attained by minimizing the following.

Communications Overhead — Macropayment schemes often establish a real-time connection to a third party for authorization during payment. In mobile billing such an online connection is established to the HLR as part of user authentication. A micropayment should be *offline* to reduce communications time and cost. This increases the possibility of fraud; hence, the cost of committing such fraud should be made greater than the possible gain.

Computations Performed — In order to verify a payment efficiently, the number of computationally expensive operations needs to be minimized. Asymmetric (public key) cryp-

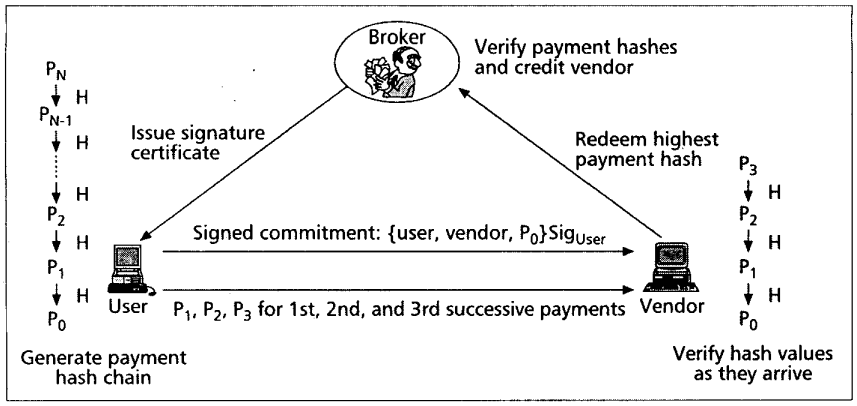

■ **Figure 4.** *Micropayments using hash chains.*

tography, such as the RSA algorithm, is more compute-intensive than symmetric cryptography, which in turn requires more computation than hash functions such as MD5 and SHA [23]. The exact speed differences will depend on the algorithms and implementations. Typically, a hashing operation will operate four orders of magnitude faster than RSA signature generation and three orders of magnitude faster than RSA signature verification. Micropayment systems will maximize the use of computationally fast operations, such as hash functions, while eliminating any public key operations where possible. This efficiency allows thousands of payments to be processed per second.

Micropayment research has concentrated on repeated payments to a *single* vendor, as in the ASPECT scheme. Many of these systems are based on the use of one-way hash functions to generate chains of hash values. Lamport originally used such hash chains for access control [24]. Pederson's "phone ticks" later used them to pay a single network operator for a phone call [25], as part of the ESPRIT project CAFE [26]. The ASPECT payment scheme was also based on phone ticks. Further schemes which apply the use of hash chains to encode amounts for payment include PayWord [27], Netcard [28], iKP micropayments [29], and PayTree [30]. The basic idea of these schemes is that a user generates a *hash chain* by repeatedly applying a hash function to a random value P_N. The user *commits* to the hash chain by digitally signing a message containing the final hash value P_0, and sends it to the vendor. For each micropayment the user releases the next payment hash, the *pre-image* of the current value, to the vendor. For the first payment P_1 is released, for the second payment P_2, and so on. Since the hash function is one-way, only the user could have generated this value, and knowledge of it can constitute proof of payment. Actual monetary value is claimed by redeeming the spent hash tokens, along with the commitment, at a broker with whom the user has an account. The process is illustrated in Fig. 4.

Multiparty Payment in Real Time

With mobile communications a user might arrive in a new network, place calls which are routed through several independent networks, and use the services of both local and remote VASPs. Consider the scenario illustrated in Fig. 5. Upon arrival in a new city, a mobile visitor places a call through the local mobile network operator to VASP1 to obtain a city traffic report and directions to his hotel. He then calls an acquaintance, another mobile user located in the same network, to inform her of his arrival. Finally, he places a long distance call, which is routed through two independent networks, to the remote VASP3 who provides him with voicemail services. We propose two different solutions, each employing micropayment technology, which allow all entities involved in these calls to be paid as they provide the services.

A call can be imagined as a linear connection of N entities. In the first solution the mobile pays the aggregated call cost from all entities involved in the call to the local network operator (NO). That NO subtracts the amount it is due and pays the remainder to the next entity downstream in the connection. The NO that receives this payment does the same for its downstream entity and so on until the final entity gets paid. The largest payment will be from the mobile to NO1, with decreasing amounts being paid further down the line. For example, consider a call involving two NOs. The mobile might pay NO1 5 cents/unit time, of which 3 cents might be paid to NO2. Thus, NO1 charges 2 cents and NO2 charges 3 cents per unit time for providing their part of the call.

In the second solution the mobile uses the same payment token to pay all the entities at once. The mobile sends a payment token to the local NO, who forwards a copy to all the downstream entities. The payment token is worth a different amount to each entity, and this amount is fixed at call setup. Thus, NO1 could redeem the token for 2 cents, while NO2 could redeem the same token for 3 cents. We examine the details of each scheme in the following sections.

Protocol Goals

Both existing and proposed mobile billing systems have been outlined, and a vision of the architecture of future mobile networks has been presented. We now discuss the design goals of the multiparty payment solutions for these mobile networks. Particular attention is given to features which improve on existing CDR billing and solve the problems these systems face in a large multi-operator environment.

Real-Time Payment Anywhere — A mobile user should be able to pay all parties involved in a call in real time, regardless of his current location and without need for online contact with a distant HLR. By removing the need for subscriber billing, all of the associated costs imposed on network operators are eliminated. Existing mobile systems use strong authentication of users or equipment through a possibly distant home location. The purpose of this is for billing, location management for incoming calls, and key management for ciphering. Once the user pays in real time, the home location does not need to be involved. In regard to ciphering, encryption key establishment can take place using the service provider's public key. The need for subscription with a home operator can be completely removed if location management is treated as one of many services provided by VASPs.

Remove User Trust and Accountability — The number of mobile users is far greater than the number of VASPs, which in turn are more numerous than NOs. Therefore, the users should be the least trusted entities within the system. Existing credit-based mobile billing systems trust the user to pay his/her bill, based on initial strong identity verification, credit history checks, and strong online authentication at the start of a session. Unlimited credit with post-fact punishment is too open to abuse in a large global system. Extensive blacklists of stolen identities and equipment must be maintained to curb fraud and credit abuse. With so many mobile users it is desirable to remove the need to trust them, and thereby minimize fraud.

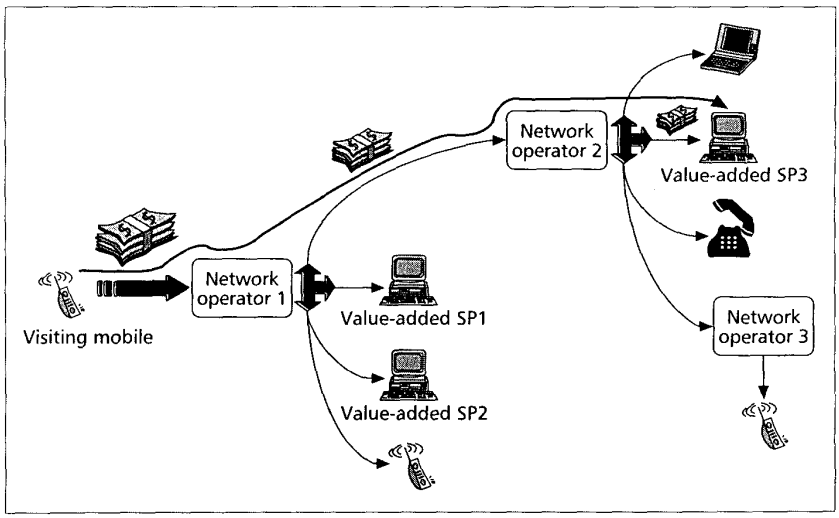

■ **Figure 5.** *Multiparty real-time payment.*

No User Signatures and Certificates — Mobile users are the least trusted entities in the system due to their numbers. While one can have some faith in a digitally signed document from an NO, a digital signature from a random roaming user, of which there can be many millions, is of less value. Use of signatures implies the existence of a public key infrastructure (PKI) capable of handling several hundred million certificates, assuming only one certificate per user. This is a huge task, especially considering that certificates will need to be revoked and the validity of a certificate checked by each party wishing to verify a user's digital signature. In a global scenario, where the services of a very large number of independent entities can be used, the use of user digital signatures, with revocation checks, as a guarantee of payment will not be efficient or scalable. Also, current mobile devices, such as GSM SIMs, are not yet capable of generating public key digital signatures. In order to improve scalability and remove the need to trust the user for payment, user digital signatures and certificates are not used within the system. This has the added advantage of affording the user more privacy.

Prevent Interoperator Fraud — Current CDR billing is based on trust and does not provide nonrepudiation. A network operator can forge CDRs. A mobile user can also deny making a call and hence refuse to pay the bill. With a large number of NOs and VASPs the possibility of any fraud between these entities needs to be removed.

Minimize Roaming Agreements — Currently operators must have roaming agreements with a foreign network in order for one of their subscribers to be able to roam into and place calls from that network. This can result in a large number of bilateral roaming agreements. Such agreements should not be necessary to allow mobile roamers to make calls in whatever network they find themselves using.

Dynamic Charging — Both NOs and VASPs should be able to dynamically price their services on a per-call basis. Tariffs can then be adjusted depending on current network conditions and quality requested, among other factors. Current VASPs are restricted by the NO's pricing model. Dynamic tariffs will allow a larger variety of services to be provided with different charging models.

Payment Flexibility — Current payphone solutions require the appropriate coins or prepaid cards to be able to pay the local NO for the call. Our second solution overcomes this by allowing a call to be paid for using tokens specific to any entity that appears in the call route.

Offline Payment Verification — Any entity accepting payment should be able to efficiently verify its validity offline, without need to contact a third party. Each payee should be guaranteed to be able to redeem a valid token with a broker.

Multiple Brokers — A payment token should be redeemable and verifiable at any broker who trusts the issuing broker. It should not be possible to spend the same payment token twice or for a single payee to redeem the same token with more than a single broker.

Identified Payee(s) — A payment token should only be redeemable by specific identified payee(s). This is to prevent tokens being stolen by eavesdroppers and cheating entities.

Token Portability — The tokens can be stored on a smart card. This allows them to be used in any mobile (or fixed) device either owned by the user, rented, or shared. A multiapplication smart card would allow a macropayment instrument such as electronic cash or a credit card scheme to be carried on the same card and used to buy payment chains online from a broker.

In summary, we wish to remove unnecessary trust from the system, reduce the online communications overhead of contacting a home location, and allow real-time payment anywhere by anyone who holds valid payment tokens.

Solution 1: Each Party Pays the Service Provider Downstream

Payment Chain Purchase from a Broker — A mobile user buys prepaid value, through their phone or terminal, from an online broker. The purpose of the broker is to aggregate micropayments between entities. To facilitate the purchase we envisage the broker being reachable through a toll-free number. A macropayment scheme such as a credit card or electronic cash is used to make the purchase. The payment chain purchase protocol is shown in Fig. 6. The mobile user creates the actual *payment hash chain* by repeatedly applying a one-way hash function, such as the Secure Hash Algorithm (SHA), to a root value P_X. The payment chain will be spendable at a specific service provider, called the *enforcer*, nominated by the user. The chain has no monetary value until committed to by a broker. To obtain this commitment the mobile user makes a macropayment to the broker, sending along the final hash (P_0), the chain length (X), the desired total value of the chain, and the identity of the enforcer through whom it must be spent, all encrypted with the broker's public key. It is assumed that the user has securely obtained and verified the broker's public key certificate beforehand. The root hash (P_X) from which the rest of the chain can be generated never leaves the mobile during the chain purchase phase.

The broker commits to the hash chain, or promises to honor its value, by digitally signing the *payment chain commit-*

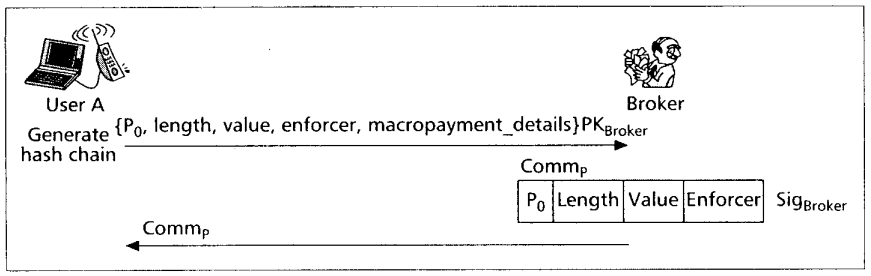

■ **Figure 6.** *A payment chain purchase.*

ment ($Comm_P$), consisting of the chain details sent by the mobile user. The commitment shows that each *payment hash value* from the chain represents a prepaid value redeemable at the broker. Each payment hash is worth the same amount, that is, the total chain value divided by the chain length. The commitment is returned to the user. If the mobile device is not capable of public key cryptography, a shared symmetric session key can be used to protect the macropayment details over the air interface, as with GSM devices. The signature on the commitment is verified by the payees at call setup, and will be rejected if invalid. This offers some security to a mobile device with no public key signature verification capabilities.

Payment hashes are released sequentially to the service provider as payment throughout a call. By fixing the enforcer in the commitment, the mobile cannot spend payment hashes more than once by attempting to *double spend* at other providers.

Assembling a Pricing Contract — To place a call the user sends the call details, such as destination, call type, QoS requirements, and payment chain commitment to the nominated service provider he is about to use. A signed *pricing contract* is then generated by the service providers involved in the call. The pricing contract has two purposes. First, it allows dynamic tariffs and different charging schemes which can be verified by the mobile user. Second, it is used to exchange payment chain commitments and to fix the starting hash for the call.

A call tariff may vary according to the service requested, current network load, and time of day, among other things. The pricing contract describes the tariff rate and charging mechanism (e.g., per second or data unit, QoS type) each SP will apply for providing their part of the service requested by the user. Since the pricing contract is signed by each SP, the mobile user can be assured that the tariffs are genuine and have not been inflated by the enforcer SP.

Figure 7 shows the three-way handshake protocol used to construct the pricing contract. In the first step each SP involved adds a line to the pricing contract. Each line consists of the following fields, as shown in Fig. 8:

- Transaction identifier, locally unique to the SP. By combining the local transaction IDs and the SP identifiers, a unique identifier for the contract is obtained.
- Service provider identity.
- Charging mechanism (e.g., duration, volume of data) and corresponding tariff rate for the SP.
- Payment chain commitment, signed by a broker, spendable only at this SP, the enforcer.
- Starting payment hash. This will be different from the hash value in the commitment if some of the chain has already been spent.
- Index position of the starting payment hash in the chain.

Each party provides the downstream entity with the payment commitment ($Comm_X$) to be used to pay them, the starting payment hash (P_{StartX}) and its position in the chain, along with the partially constructed contract.

Each entity will use a *different payment chain* to pay the downstream entity. For example, in Fig. 9 the mobile user pays SP1 using payment hash chain P, while SP1 pays SP2 using payment chain Q, and SP2 pays SP3 using payment chain R. Once constructed, the full contract is digitally signed by the service providers to prevent any fields being changed.

By including the commitments in the contract, spent payment hashes can be associated with a specific call. This creates a record of the call, providing much of the same information obtainable from a traditional CDR. When a pricing contract is associated with spent payment hashes the call duration, call route, and call destination can be obtained. Like regular CDRs, this will allow it to be used for many secondary functions such as intrusion detection, network planning, marketing data collection, and law enforcement. If desirable and depending on the macropayment system used, the identity of a mobile user can be associated with a payment chain commitment when purchased at the broker. This would allow the caller's identity to also be associated with the call. If user anonymity is required, the payment chain must be purchased using an anonymous macropayment system like some forms of electronic cash.

In step two each SP digitally signs the fully assembled pricing contract, checking that their line has not been altered in any way. The signing starts with the final SP in the call route, who passes the partially signed contract back along the route to SP1. In an implementation each SP will sign a hash of the contract. The signatures prove that each SP took part in the call and is due payment.

After SP1 has signed, the finished contract is forwarded to each SP in step three. There is no need for SPs to trust each other since each can independently verify the fully signed con-

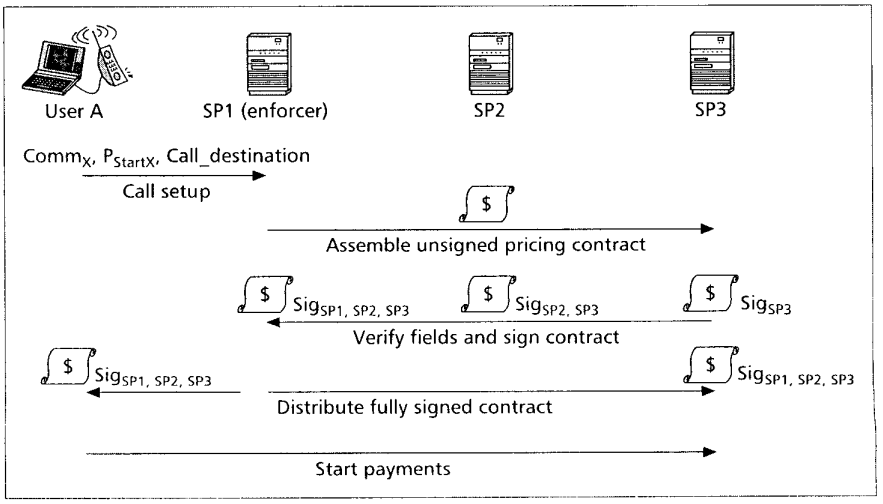

■ **Figure 7.** *Constructing a pricing contract.*

Trans1	SP1	Charge1	$Comm_1$	P_{Start1}	Start1
Trans2	SP2	Charge2	$Comm_2$	P_{Start2}	Start2
⋮					⋮
TransN	SP*N*	Charge*N*	$Comm_N$	P_{StartN}	Start*N*

$Sig_{SP1, SP2, ... SPN}$

Figure 8. *Signed pricing contract for solution 1.*

tract. The pricing contract is presented to the user for agreement before the call is set up. From the charging information fields the total call cost per charging unit is obtained. The user can verify the signatures to prove that each quote is genuine.

If necessary, the public key certificates for the service providers are distributed to each other and the mobile in steps two and three. A service provider's certificate is required in order to verify a digital signature from that service provider. In a mobile device with limited computational capability, the signature verification can be omitted as long as the user is prepared to pay the price quoted. To detect overcharging by SP1, the validity of the pricing contract can be checked later on another device or with a broker. A new contract may be established midcall to reflect any changes in tariff. For example, this might occur if a long call overlaps the switch from peak rate to off-peak rates. Similarly, if an interoperator handover occurs during a call, a new pricing contract is established as part of the call setup with the new operator.

Making Payments — The user pays the total charge from all service providers to the first service provider in the call. That service provider subtracts the amount it is owed and in turn pays the downstream entity the remaining amount. Having agreed to the pricing contract, the user begins the call by releasing a payment hash worth the total amount due per charging unit. For example, in Fig. 9 SP1, SP2, and SP3 might charge 2, 2, and 1 cent(s) per unit, respectively, for the call, yielding a total charge of 5 cents. If $Comm_P$ represents a new payment chain with hash values worth 1 cent, P_5, the fifth hash in the chain is sent to SP1. SP1 applies the hash function five times repeatedly to this payment hash, and if the payment is valid the result will be the final hash P_0. The value of a payment hash is fixed in the broker commitment. Sending P_5 after P_0 is equivalent to sending all five payment hashes, P_1–P_5, since they can be obtained from P_5 by repeatedly hashing. In turn, SP1 pays SP2 using payment hashes worth 1 cent from the chain defined by $Comm_Q$. SP1 pays SP2 the remaining 3 cents by sending Q_3,. SP2 applies three repeated hash functions to Q_3 to obtain Q_0, proving the payment is valid.

The flow of payment continues downstream until the last SP is paid. Payment is *ongoing*, with the user releasing hashes at regular intervals according to the charging mechanism. For a voice call this might be every second. In return for a valid payment, the SPs continue to provide the service they agreed to in the pricing contract. If the user does not receive these services, he can terminate the call by not releasing any more payment hashes. In Fig. 9, *N* is the number of payments made. When paying another SP downstream, the payer may use payment hashes from a different broker than the upstream entity.

Redeeming Tokens — At the end of the day the SP sends the pricing contract, the highest spent payment hash, and the position of that hash from the final hash (P_0) in the chain to the broker:

$$\{Contract, P_Y, Y\}Sig_{SP}$$

A broker will only redeem payment from the SP identified in the commitment, with a matching pricing contract. This prevents unauthorized parties from redeeming stolen payment hashes.

The broker verifies that the payment is valid by performing *Y* hashes on P_Y and comparing the result with P_0 in the commitment. He validates his signature on the commitment and checks his records to see that this part of the chain has not already been redeemed by the SP. To limit the state that must be held, certificates can expire after several months, with payment chains only being valid while the certificate is valid.

After verifying that the chain is valid, the broker pays the service provider the total amount redeemed. Any *unspent payment hashes* from the chain can be spent later through the same SP. It is up to the SP to remember the last spent hash from the chain, to prevent double spending of used payment hashes. A new pricing contract is established with the original commitment and the last spent hash as the new starting one. Figure 9 shows that SP2 starts to pay SP3 with an unspent hash R_{21} from a partly used chain with commitment $Comm_R$. The user and SP1 start with hashes P_5 and Q_3, respectively, both from new chains.

A later section describes how an extra field can be added to the pricing contract so that payment hashes may be redeemed indirectly through any broker. A mobile user might roam out of an NO's coverage area before all the value in a chain has been spent. To allow unspent value from a payment chain to be claimed back from the broker, the chain can be given an expiry date. Payment hashes which have not been spent or redeemed before this time can be exchanged for new value by the user upon presentation of the commitment and highest payment hash.

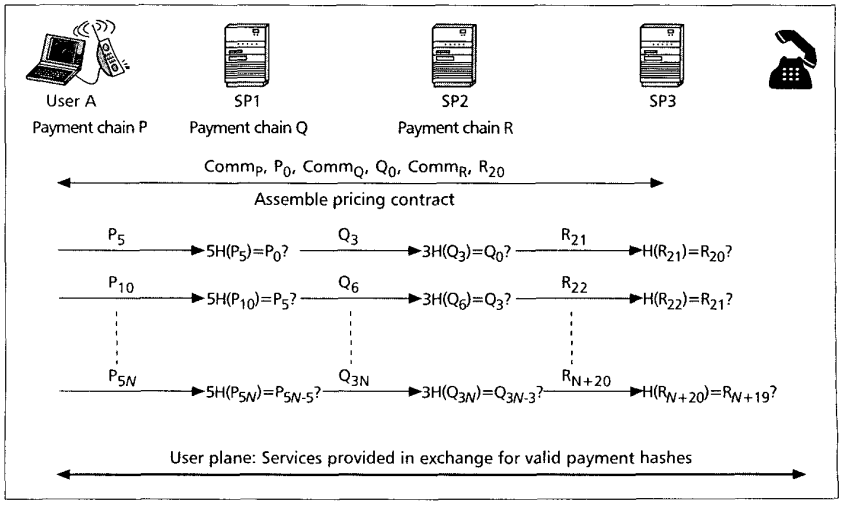

Figure 9. *Solution 1: Each party pays the entity downstream.*

Discussion — The first solution applies payment hashes on a link-to-link basis with each entity paying the downstream entity the total remaining call cost. A pricing contract is introduced to allow verifiable dynamic tariffs, fix the starting hash for each payment commitment and to create an undeniable record of the call.

These advantages come at an additional computational cost per call over traditional CDRs. The cost is:

- 1 signature/SP at call setup for the pricing contract
- 1 signature verification/SP on the payment commitment
- $N^*(A/V)$ hash functions/SP to verify payment, where A is the amount due per charging unit, V is the value per payment hash in the chain being used, and N is the number of payments made
- x signature verifications on the pricing contract by the mobile and optionally each SP, where x is the number of SPs in the call
- 1 database lookup/SP if a partly spent chain is to be used, to ensure that the starting hash is equal to the highest spent hash from the chain

However, of these only the computationally efficient hashing functions are used during the call. The mobile, broker, and SPs will also have additional offline computational costs of buying and redeeming payment chains. The computational overhead is far less than in existing macropayment systems. The computational and communications cost of online authentication through a remote home location is removed. Existing CPU speeds and cryptographic accelerators are capable of handling the load for thousands of simultaneous calls. We believe the advantages gained are worth this additional computational cost.

The solution achieves all the protocol goals described earlier, except payment flexibility. In particular, the advantages over traditional CDRs in a mobile environment are the removal of trust between the user and SP, and also between the SPs themselves. Each SP is guaranteed payment, and the user is guaranteed to only have to pay the tariffs presented at call setup. With traditional mobile billing the exact cost of the call is not presented to the user beforehand. The need for user authentication, online contact with a home location, and roaming agreements is eliminated. Additionally, the functions of traditional CDRs, other than billing, are not lost. The broker and SPs can independently obtain a record of the call from the pricing contract and highest spent payment hash.

Solution 2: Mobile Pays All SPs Directly

The main disadvantage of the first solution is that the local cellular NO through whom the call is to be placed must be known in advance. If the mobile user roams away from this NO he will have to buy a new payment chain for use with the new NO.

Our second solution solves this problem by allowing the *same payment hash* to be spent at all SPs participating in the call at the same time. When the user purchases a payment chain, she must still nominate an enforcer through whom to spend it, but this can be *any* network operator or VASP in the call route. For example, a user might pay a VASP to provide her with voicemail facilities. She buys a payment chain to spend at the voicemail provider. Now, whenever she roams and whichever networks she uses to access her voicemail, she can use that chain to pay not only the voicemail VASP but also all the NOs through which she is connecting.

If a call is being made to a party other than the voicemail provider, the same payment chain should not be used. Instead, a different payment chain specific to one of the SPs involved in the call route will be employed. This prevents a geographically distant SP, not found in the route of the current call, being unnecessarily involved. A mobile user might typically carry several payment chains at once. Each chain would be specific to an SP through whom or to whom she frequently places calls. As long as the SP is present in the route of the active call, the chain can be used to pay all SPs involved in that call.

The protocol used to purchase a broker-signed commitment to a user generated payment chain is identical to that used in the earlier solution. However, in the first solution the value of each payment hash in the chain was fixed to be equal to the total chain value divided by the chain length. In the second solution the monetary value of a single payment hash is not fixed, allowing the same hash to be used to pay all parties, without the possibility of fraud. The enforcer will prevent more than the total value of the chain from being spent. Failure to do so will be detected by the broker when the hashes are redeemed. The chain length is still included in the commitment, so the enforcer does not set the hash value such that it requires more hashes than are in the chain to be used to spend the total value.

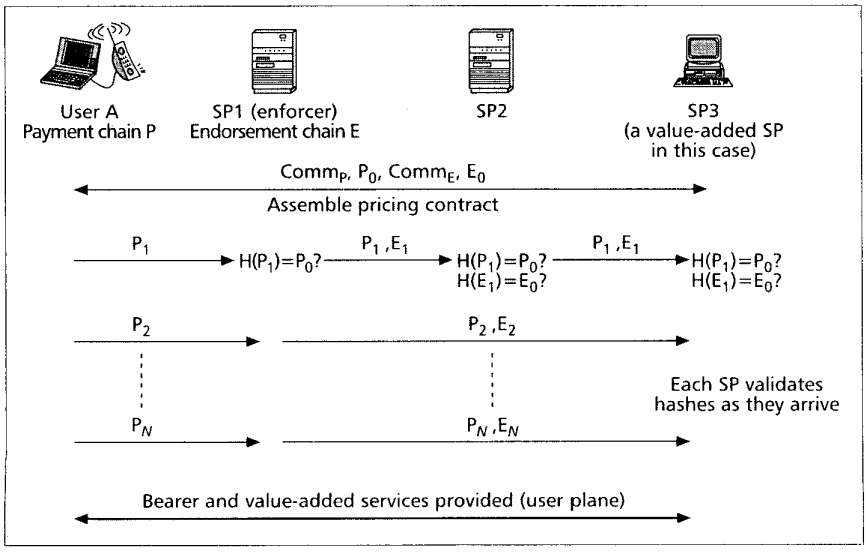

■ **Figure 10.** *Solution 2: The mobile pays all SPs with the same payment hash, with SP1 as the enforcer.*

Figures 10 and 11 show the same call being made, but with payment chains for different enforcer SPs along the route. Figure 10 shows how the scheme works when the local NO is the enforcer. In Fig. 11 the payment chain must be spent through the VASP, SP3. This could be the voicemail provider in our earlier example. The details of the payment mechanism, including the construction of a pricing contract and function of the endorsement chain, are now examined.

Pricing Contract — In the second solution the pricing contract has more significance. In addition to providing the original functions, it is also used to link a single payment commitment to multiple SPs for a call. A problem arises when a pay-

ment hash may be redeemed by any party, as in the second solution. Parties who were not involved in a call and are not entitled to payment may try to redeem payment hashes. The pricing contract is used to identify those parties who may redeem payment hashes from a specific chain. An SP can only redeem a payment hash from a broker if it has a valid pricing contract which authorizes it to do so.

The pricing contract has a different structure in the second solution, as shown in Fig. 12. It consists of:

- The combined local transaction IDs from each SP.
- The identity of each NO and VASP involved in the call. When combined with the transaction ID, a unique identifier for the contract is obtained.
- Charging mechanism and tariff rate for each SP.
- Payment chain commitment: a single payment chain, from the mobile user and spendable through the enforcer.
- Starting payment hash from the chain for the current call.
- Position of that hash in the chain.
- The value per payment hash for the duration of the call.
- Endorsement chain commitment, signed by the enforcer SP. This prevents double spending of payment hashes by the mobile user or SPs.

The important changes to the pricing contract are that only a single payment chain is present, the value of hashes from that chain are fixed in the contract, and an endorsement chain is present to prevent double spending of payment hashes.

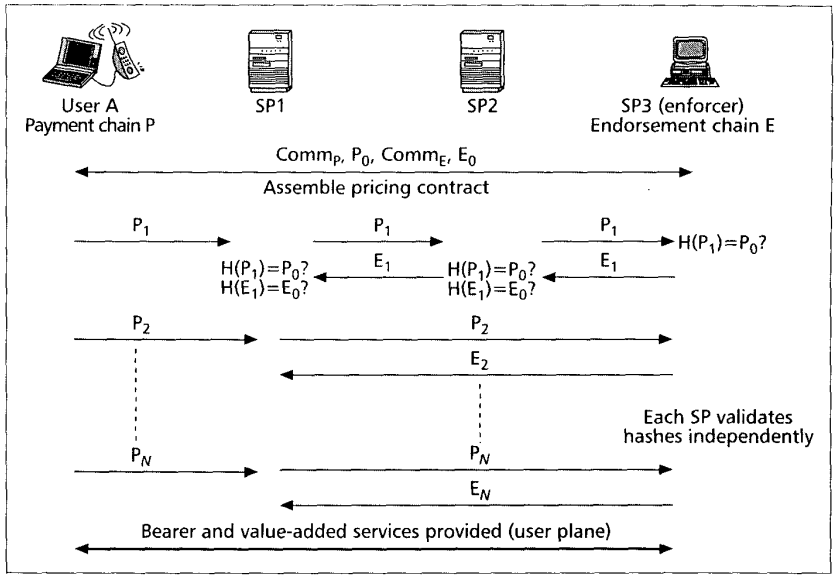

■ **Figure 11.** *Solution 2: The mobile pays all SPs with the same payment hash, with SP3 as the enforcer.*

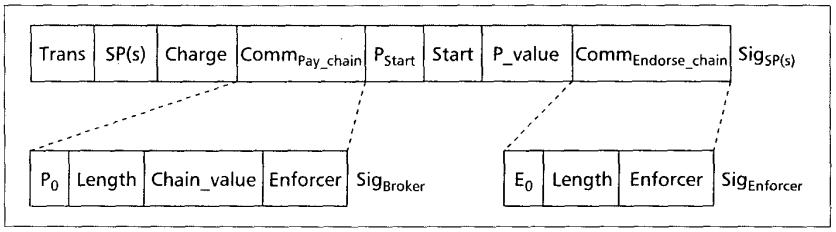

■ **Figure 12.** *Contents of a pricing contract, payment, and endorsement commitments.*

The enforcer is responsible for ensuring that it is constructed correctly using the same three-way handshake protocol described earlier. As before, each SP signs the contract to prove that they took part in the call and are due payment. The enforcer signs the contract last. This prevents a partial contract being replayed to trick another SP into accepting an old contract with an already spent hash chain. The enforcer knows which payment hashes have been spent, since they must pass through it, and it will not sign a contract with an already spent starting hash.

In Fig. 10 SP1 signs the contract last, since it is the enforcer. However, when the pricing contract is being assembled for the scenario in Fig. 11, the three-way handshake protocol will occur in the opposite direction, allowing the enforcer SP3 to sign the contract last.

The total call cost per unit time, or per data unit transmitted, is the sum of each SP's tariff rate in the pricing contract. Assume for the moment that all SPs use the same unit of measurement for which they must be paid. For example, this might be per second for a voice call or per kilobyte sent for a data call. Each payment hash is worth the total cost per charging unit. This is best illustrated by example. Three SPs provide voice bearer services for a call charging 0.1, 0.5, and 0.2 cent/s, respectively. The enforcer assigns each payment hash to be worth 0.8 cent in the pricing contract. Only a single hash needs to be released every second, and this is redeemable by each SP. When the service providers redeem P_{40}, the 40th payment hash, they will be paid 4, 20, and 8 cents, respectively. The broker knows how much to pay each SP from the contents of the pricing contract. If the broker cannot trust the enforcer to assign value to payment hashes, it can be fixed in the broker commitment, as with the first solution.

Endorsement Chains, Double Spending, and Change — The enforcer, identified in the payment chain commitment, is given the role of preventing double spending of payment hashes. Since all payment hashes must pass through the enforcer, it can keep a record of how much of the chain has already been spent.

After a call finishes, the mobile can reuse unspent hashes (the change) on another call which may pass through different networks to a different destination than the previous call. During the call, the enforcer will ensure that the payment hashes it and the other SPs receive have not already been spent. However, without further protection mechanisms, cheating can take place after this call. For example, the user could give hashes spent in the second call to a different SP from the first call. Now both SPs from both calls have valid pricing contracts for the hashes, which they can redeem from the broker. Similarly, SPs from each call could collude to swap payment hashes to gain value in this way. While the broker can detect this fraud when the same hash is later redeemed twice, he cannot be sure of who committed the fraud and thus who should not get paid.

To solve this problem we introduce the concept of an *endorsement chain*. This is a hash chain created and committed to by the enforcer for each call. It consists of a final hash (E_0), chain length, and enforcer ID, all signed by the enforcer:

$$Comm_{Endorse\_chain} = \{E_0, Length, Enforcer\}Sig_{Enforcer}$$

There is no value associated with an endorsement chain; its sole purpose is to prevent double spending. The enforcer's new endorsement commitment is included in each new pricing contract constructed. To make a payment the mobile user releases a single payment hash per charging unit. The enforcer applies a single hash function to verify it, and compares the result to the last received hash. If valid, the enforcer attaches a corresponding endorsement hash to each payment hash before forwarding it to the other SPs, as shown in Figs 10 and 11. This endorsement hash indicates that the enforcer SP accepted the corresponding payment hash. An endorsement hash is specific to a call described by the pricing contract containing the endorsement and payment chain commitments.

To validate a payment each SP must verify the payment and endorsement hashes by recomputing the hash function on them. If both are valid, they will hash to the previous values received, or the commitment values in the case of the first payment.

The broker will now only accept a payment hash from an identified SP if a corresponding endorsement hash and pricing contract accompany it. In this way, double spending by the user and SPs other than the enforcer is prevented. As before, only the highest hash from both chains need be sent to the broker along with the pricing contract:

$$Redeem = \{Pricing\_contract, P_X, X, E_Y, Y\}Sig_{SP},$$

where X and Y are the positions of the payment hash and endorsement hash, respectively in the hash chains specified by the commitments in the pricing contract. Double spending by the enforcer SP cannot be prevented since it is entrusted with generating the endorsement hashes. However, if the enforcer does cheat, it will be detected after the fact when other SPs redeem the same payment hashes twice. Post-fact detection is acceptable for the enforcer because the broker will refuse to issue payment chains in the name of enforcers it does not trust.

Broker Clearing — One disadvantage of the schemes, especially in the second solution, is the requirement that SPs redeem payment hashes from the issuing broker. For geographically dispersed SPs this will introduce a communication overhead, even when performed offline. To address this limitation, a network of brokers may be used whereby a payment chain may be redeemed at any broker agreed upon at the time of call setup. When the pricing contract is constructed each SP fixes the *redeeming broker*, normally a local broker, with whom he is going to redeem the payment chain. This is an extra field per SP in the pricing contract:

$$\{SP1: BrokerA, SP2: BrokerB, \ldots SPN: BrokerX\}Sig_{SP1,SP2..SPN}$$

No other broker can now redeem the part of the chain spent during the call; hence, double redeeming is prevented. The redeeming broker later clears payment chains in bulk with the issuing broker. Existing financial clearing networks could be enhanced to exchange these details. One can envision a broker per area or region who will redeem for multiple SPs in that area.

Fraud is only possible when the enforcer signs multiple pricing contracts for the same part of a payment chain with different redeeming brokers for an SP. The SP can then redeem the same chain from different brokers. If an enforcer cannot be trusted in this case, the issuing broker can nominate, in the payment commitment, a small number of brokers or even a single broker, as the redeeming brokers.

Discussion and Critique

Both solutions provide an efficient means of allowing real-time payment to multiple service providers with dynamic tariffing and variable charging schemes. Incontestable charging and guaranteed payment are provided without the need for user authentication or online contact with a remote home location. A micropayment scheme using hash functions and offline broker contact allows the solutions to be efficient and scalable. To aid performance, digital signatures are only used at call setup to generate the pricing contract. Unlike traditional billing, the schemes allow fraud prevention to be decoupled from the clearing process.

The main disadvantage of each solution is that a service provider which will be involved in the call must be known in advance. In the first solution this must be the local SP, while the second solution is more flexible, allowing it to be any SP. However, payment chains are purchased online from a broker at any time, and this process can be seamlessly integrated with call setup. In addition, the requirement to use a specific SP is no different than current prepaid phone cards, calling cards, or discount calling services.

When comparing both solutions, the first offers the service providers more flexibility. A long payment chain between two SPs can be used for many calls. Where two SPs trust each other, credit hashes can be used instead of prepaid ones. By establishing a pricing contract just between two SPs, a payment chain may be used to pay for multiple simultaneous calls between them. Solution one is also suitable for use in a migration scenario, where some SPs still use traditional CDR billing. Instead of being paid or paying with hashes, a legacy SP will generate CDRs and bill other SPs.

The second solution allows the value of a payment hash to be fixed at call setup, and the same hash to be used to pay all SPs in the call. An enforcer is used to prevent double spending by issuing an endorsement hash for each payment hash. The computations necessary in solution two are the same as in the first solution, except for:

- Two signature verifications per SP, one to verify the payment commitment, the other to verify the endorsement commitment.
- $2N$ hash functions per SP, where N is the number of payments made; one hash function for the payment, one for the endorsement.
- Additional signature by the enforcer on the endorsement chain. Endorsement chains can be generated in bulk at the beginning of the day.

Thus, the number of hash functions performed during the call is lower, but each SP needs to perform an additional signature verification on the endorsement chain at call setup. For calls involving more than two SPs the size of the pricing contract will be smaller, due to fewer commitments being present.

Solution two provides more flexibility to the user by allowing any SP in the call to act as an enforcer to a payment chain. When a mobile user roams from one network to another, it is not necessary to purchase a new chain, provided the enforcer is present in the call route. The same payment message from the enforcer is passed along through all other SPs, requiring no new message construction.

Enhancements and Hybrids — Situations may arise where SPs in a call use different charging mechanisms. For example, a VASP might charge based on data content, whereas intermediate NOs might charge on the volume of data transmitted. In such a case a second payment chain can be used to pay the VASP, with the appropriate additional information being added to the pricing contract.

We have applied the solutions to reverse charging, split charging, and group charging scenarios. Split charging refers to when the call cost is divided between two parties, such as

when calling a roaming mobile user. A group call involves three or more end parties in a branching network topology. While we found that the construction and negotiation of pricing contracts required extra effort, the actual payment method remained largely unchanged in each case.

User-to-user payments are also possible provided the payee has a certificate, allowing him/her to redeem payment hashes. While special hardware is not required in our scheme, a smart card can be used to securely store and transport payment chains if desired.

Implementation Status

A prototype has been developed in Java using cryptographic libraries to supply certificate functionality, the RSA algorithm for digital signatures, and MD5 and SHA hash functions. Java Remote Method Invocation (RMI) was used to handle message passing across the network. NOMAD [31], an application based on a popular Internet telephony package, was used to demonstrate payment with personal mobility for a voice call through multiple service providers.

Conclusion

With a large number of network operators and value-added providers, it is necessary to guarantee payment and remove the complex trust relationships involved in billing. Traditional billing methods for mobile systems, based on the generation of a CDR, are examined. Based on these a number of problems are identified, especially in relation to future mobile networks with many service providers. The desirable properties of a mobile payment system are drawn up, and two solutions which securely achieve these goals proposed. The additional computational cost of each solution over traditional CDR generation was presented. We believe that the small additional computation is an acceptable cost for the features, listed earlier, gained over CDR billing.

To allow an efficient solution, with minimum computational cost during a call, a scheme based on hash chains was used. Micropayment research has concentrated on providing payment to a single vendor at any one time. We have proposed a solution which allows micropayments to multiple vendors at the same time.

Dynamic tariffs and different charging schemes are possible by constructing a pricing contract at call setup. We have eliminated the need for user certificates and authentication. This greatly reduces the number of public key certificates needed in the system, and reduces the computational load for the user. In addition, it removes the need for contact with an HLR and provides a desirable level of anonymity and privacy. Migration using a mixture of payment hashes and CDRs is possible. The solutions are also suitable for call services requiring split and group payment.

Our solutions provide an efficient means of ensuring real-time payment in a multi-service-provider environment without the need for user authentication. We expect the challenge of mobile payment methods to grow dramatically in importance as mobile communications become increasingly sophisticated and ubiquitous.

Acknowledgments

The authors would like to thank the anonymous reviewers for their valuable comments and suggestions.

References

[1] J. Flood, *Telecommunication Networks*, 2nd ed., London: IEE, 1997.
[2] F. Redmill and A. Valdar, *SPC Digital Telephone Exchanges*, London: IEE, 1994.
[3] ETSI GSM 12.05, "Event and Call Data," Aug. 1998.
[4] GSM Assn. PRD TD.17, "Data Record Format: Version 2," Apr. 1998.
[5] GR-1100-CORE, "Bellcore Automatic Message Accounting Format Generic Requirements," Bellcore, 1998.
[6] ETSI I-ETS 300 819, "TMN; Functional Specification of the Usage Metering Information Management on the Operations System/Network Element Interface," June 1998.
[7] ITU Q.825, "Draft Specification of TMN Applications at the Q3 Interface: Call Detail Recording," 1998.
[8] ETSI GSM 01.02, "General Description of a GSM Public Land Mobile Network (PLMN)," Mar. 1996.
[9] GSM Assn., http://www.gsmworld.com
[10] GSM Assn. PRD BA.07, "Charging and Accounting Principles," Aug. 1998.
[11] GSM Assn. PRD BA.12, "Transferred Account Procedure and Billing Information," Aug. 1998.
[12] CIBERNET, "Cellular Industry Billing Exchange Record (CIBER)," Washington, DC, 1997.
[13] Y. Cho, Y. Lin, and H. Rao, "Reducing the Network Cost of Call Delivery to GSM Roamers," *IEEE Network*, vol. 11, no. 5, Sept. 1997, pp. 19–25.
[14] ITU-T Recommendation D.140, "Accounting Rate Principles for International Telephone Services," July 1998.
[15] Ericsson Prepaid Solutions, http://www.ericsson.com/systems/gsm
[16] D. O'Mahony, "UMTS: The Fusion of Fixed and Mobile Networking," *IEEE Internet Comp.*, vol. 2, no. 1, Jan. 1998, pp. 49–56.
[17] ACTS Project AC095 ASPECT, http://www.infowin.org
[18] G. Horn and B. Preneel, "Authentication and Payment in Future Mobile Systems," *Comp. Sec. — ESORICS '98*, LNCS 1485, Berlin: Springer-Verlag, 1998, pp. 277–93.
[19] K. Martin *et al.*, "Secure Billing for Mobile Information Services in UMTS," *Intelligence in Services and Networks*, IS&N '98, LNCS 1430, Berlin: Springer-Verlag, 1998, pp. 535–48.
[20] 3GPP TS 22.115, "3rd Generation Mobile Telecommunications; Service Aspects; Charging and Billing," Apr. 1999, http://www.3gpp.org
[21] GSM Assn. PRD TG.24, "Requirements for Charging, Billing, Accounting, and Tariffing," Oct. 1997.
[22] D. O'Mahony, M. Peirce, and H. Tewari, *Electronic Payment Systems*, Boston/London: Artech, 1997.
[23] B. Schneier, "One-Way Hash Functions," *Applied Cryptography*, 2nd ed., Wiley, 1996, pp. 429–59.
[24] L. Lamport, "Password Authentication with Insecure Communication," *Commun. ACM*, vol. 24, no. 11, Nov. 1981, pp. 770–72.
[25] T. Pederson, "Electronic Payments of Small Amounts," *Security Protocols*, LNCS 1189, M. Lomas, Ed., Springer-Verlag, 1997, pp. 59–68.
[26] J. Boly *et al.*, "The ESPRIT Project CAFE - High Security Digital Payment Systems," *Comp. Sec. — ESORICS '94*, LNCS 875, Berlin: Springer-Verlag, 1994, pp. 217–30.
[27] R. Rivest and A. Shamir, "PayWord and MicroMint: Two Simple Micropayment Schemes," *Security Protocols*, LNCS 1189, M. Lomas, Ed., Springer-Verlag, 1997, pp. 69–87; http://theory.lcs.mit.edu/~rivest
[28] R. Anderson, H. Manifavas, and C. Sutherland, "NetCard - A Practical Electronic Cash System," *Proc. 4th Cambridge Wksp. Sec. Protocols*, Cambridge, U.K., 1996, http://www.cl.cam.ac.uk/users/rja14
[29] R. Hauser, M. Steiner, and M. Waidner, "Micro-payments based on iKP," *Proc. 14th Worldwide Cong. Comp. and Commun. Sec. Protection*, Paris, France, 1996, pp.67-82; http://www.zurich.ibm.com
[30] C. Jutla and M. Yung, "PayTree: Amortized-Signature for Flexible MicroPayments," *Proc. 2nd USENIX Wksp. Elect. Commerce*, Oakland, CA, 1996, pp. 213–21.
[31] D. O'Mahony, L. Doyle, H. Tewari, and M. Peirce, "NOMAD — An Application to Provide UMTS Telephony Services on Fixed Terminals in COBUCO," *Proc. 3rd ACTS Mobile Commun. Summit*, Rhodes, Greece, June 1998, vol. 1, pp. 72–76.

Biographies

MICHAEL PEIRCE (Michael.Peirce@cs.tcd.ie) received his B.A. (Mod) degree in computer science and his M.A. degree from Trinity College Dublin in 1995 and 1998, respectively. He is currently a Ph.D. student in the Department of Computer Science at Trinity College Dublin, where he is a member of the Networks and Telecommunications Research Group. His research interests include network security and electronic payment for mobile communications. He is co-author of *Electronic Payment Systems*, published by Artech House.

DONAL O'MAHONY (Donal.OMahony@cs.tcd.ie) is a lecturer in computer science at Trinity College, where he coordinates a research group working in the areas of networks and telecommunications. The research group addresses ongoing projects in high-speed network technologies, multimedia collaborative applications, network security, and electronic payment. The group is currently working on the 4th Generation Telephony project, which aims to deliver universal voice connectivity in fixed and wireless environments based on an evolved IP core network. He received B.A., B.A.I., and Ph.D. degrees from Trinity College Dublin.